

Maximin equilibrium*

Mehmet ISMAIL†

March, 2014. This version: June, 2014

Abstract

We introduce a new theory of games which extends von Neumann's theory of zero-sum games to nonzero-sum games by incorporating common knowledge of individual and collective rationality of the players. Maximin equilibrium, extending Nash's value approach, is based on the evaluation of the strategic uncertainty of the whole game. We show that maximin equilibrium is invariant under strictly increasing transformations of the payoffs. Notably, every finite game possesses a maximin equilibrium in pure strategies. Considering the games in von Neumann-Morgenstern mixed extension, we demonstrate that the maximin equilibrium value is precisely the maximin (minimax) value and it coincides with the maximin strategies in two-player zero-sum games. We also show that for every Nash equilibrium that is not a maximin equilibrium there exists a maximin equilibrium that Pareto dominates it. In addition, a maximin equilibrium is never Pareto dominated by a Nash equilibrium. Finally, we discuss maximin equilibrium predictions in several games including the traveler's dilemma. JEL-Classification: C72

*I thank Jean-Jacques Herings for his feedback. I am particularly indebted to Ronald Peeters for his continuous comments and suggestions about the material in this paper. I am also thankful to the participants of the MLSE seminar at Maastricht University. Of course, any mistake is mine.

†Maastricht University. E-mail: mehmet@mehmetismail.com.

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 individually and collectively rational and letting this be common knowledge among players. In other words, we extend von Neumann’s theory of games from zero-sum games to nonzero-sum games by capturing both conflicting and cooperating preferences of the players via these rationality² assumptions.

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 nonzero-sum games. But this did not seem to be 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 to an ignorance of her opponent’s payoffs and hence the preferences of

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.

³For example, see von Neumann and Morgenstern (1944, p. 540).

	a	b	c
a	1, 1	3, 3	0, 1
b	3, 1	3, 3	3, 4
c	3, 3	0, 3	4, 0

Figure 1: A game with the same payoffs to the Nash equilibria and to the maximin strategies.

her opponent. These arguments call for a revision of the maximin strategy concept.

Let us consider the following two games to support our statement. In the first game shown in Figure 1, Alfa (he) is the row player and Beta (she) is the column player. There are four Nash equilibria $[(0, \frac{3}{4}, \frac{1}{4}), (0, \frac{1}{4}, \frac{3}{4})], [(\frac{1}{3}, \frac{2}{3}, 0), b], (a, b)$ and (c, a) . All the Nash equilibria yield the same (expected) payoff vector $(3, 3)$. On the other hand, Alfa has a unique maximin strategy b which guarantees him to receive a payoff of 3. Beta also has a unique maximin strategy b which guarantees her the same payoff of 3.

Although the point we want to make is different, it is of importance to note the historical discussion about this type of games where the Nash equilibria 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 which he calls unprofitable. Luce and Raiffa (1957) and Aumann and Maschler (1972) argue that maximin strategies seem preferable in those cases. In short, in the games similar to Figure 1, the arguments supporting maximin strategies are so strong that it led some game theory giants to prefer them over the Nash equilibria of the game. These arguments, however, may 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. Let the payoffs be as given in Figure 2 with some small $\epsilon > 0$. Notice that for every $\epsilon > 0$, the profile $(trick, trick)$ is a Nash equilibrium and that the Nash equilibria of the previous game are still Nash equilibria in this game.⁴ By contrast, notice that the maximin

⁴Depending on ϵ , the game has other Nash equilibria as well. For example when $\epsilon = 1$, the other Nash equilibria are $[(0, \frac{3}{16}, \frac{1}{16}, \frac{3}{4}), (0, \frac{1}{16}, \frac{3}{16}, \frac{3}{4})], [(\frac{1}{2}, \frac{1}{6}, 0, \frac{3}{4}), (0, \frac{1}{4}, 0, \frac{3}{4})]$ and $[(0, 0, \frac{1}{4}, \frac{3}{4}), (\frac{1}{4}, 0, 0, \frac{3}{4})]$ all of which yield 0 for both players.

	a	b	c	<i>trick</i>
a	1, 1	3, 3	0, 1	$-\epsilon, 0$
b	3, 1	3, 3	3, 4	$-\epsilon, 0$
c	3, 3	0, 3	4, 0	$-\epsilon, 0$
<i>trick</i>	0, $-\epsilon$	0, $-\epsilon$	0, $-\epsilon$	0, 0

Figure 2: A game where $\epsilon > 0$.

strategies of the previous game disappears no matter which ϵ we take.⁵

Suppose that players made an agreement (explicitly or implicitly) to play the Nash equilibrium (a, b) .⁶ Then, Alfa would make sure that Beta does 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 does not make a unilateral deviation to *trick* for the same reason. Therefore, one observes that Nash equilibria are 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. In this paper, maximin principle is included in maximin equilibrium 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 (b, b) in the second game. Suppose that players agreed (explicitly or implicitly) to play this profile. Alfa would make sure that Beta does 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 3. Beta would also make sure that Alfa does not unilaterally deviate to *trick* if he is ‘rational’. A player might still deviate to a, b or c but this is okay for both players since they both guarantee their respective payoffs in

⁵It is clear that whichever game we consider, it is possible to make maximin strategies disappear by this way.

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.

this region.⁸ That is, Alfa would guarantee to receive 3 given that Beta is ‘rational’ and Beta guarantees to receive 3 given that Alfa is ‘rational’.

In conclusion, if the value of the first game is 3 then the value of the second game should, intuitively, not be strictly less than it given the common knowledge of ‘rationality’ of the players. It seems that the profile (b, b) still guarantees the value of 3 under this assumption. We show that maximin equilibrium that we introduce in this paper captures this property.

In Section 2, we present the framework and the assumptions we use in this paper. In Section 3, we introduce a deterministic theory of games via the concept of maximin equilibrium. Maximin equilibrium extends Nash’s value approach to the whole game and evaluates the strategic uncertainty of the game by following a similar method as von Neumann’s maximin strategy notion. We show that every finite game possesses a maximin equilibrium in pure strategies. Moreover, maximin equilibrium is invariant under strictly increasing transformations of the payoff functions of the players. In Section 4, we extend the analysis to the games in von Neumann-Morgenstern mixed extension. We demonstrate that maximin equilibrium exists in mixed strategies too. Moreover, we show that a strategy profile is a maximin equilibrium if and only if it is a pair of maximin strategies in two-person zero-sum games. In particular, the maximin equilibrium value is precisely the minimax value whenever it exists. Moreover, we show that for every Nash equilibrium that is not a maximin equilibrium there exists a maximin equilibrium that Pareto dominates it. In addition, a maximin equilibrium is never 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 n -person games. All the results provided in Section 3 and in Section 4 hold in n -person games except the one which requires a 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):

⁸Note that b is the maximin strategy of both players in the 3 by 3 game.

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 sets 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 players’ actions. As usual, the notation x in $X = X_1 \times X_2$ represents a strategy profile.⁹ In short, a two-player non-cooperative game Γ can be denoted by the tuple $(\{1, 2\}, X_1, X_2, u_1, u_2)$. We distinguish between the game Γ and its von Neumann-Morgenstern mixed extension. Clearly, the mixed extension of a game requires more assumptions to be made and it will be treated separately in Section 4. When it is not clear from the context, we refer the original game as the pure game or the deterministic game to not to cause a confusion with the games in mixed extension. Starting from simple strategic decision making situations, we firstly introduce a deterministic theory of games in this section and in the following one.¹⁰

For the analysis of a game we need a notion of rationality of the players. 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 games with more than one players, however, it is not unambiguous what it means 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

⁹As is standard in game theory, we assume that what matters is the consequence of strategies (consequentialist approach) so that we can define the utility functions over the strategy profiles.

¹⁰Note that all the definitions we present can be extended in a straightforward way to n -person games which will be introduced in Section 5.

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 nonzero-sum games since the preferences of the players are not necessarily opposing. Due to the fact that rationality of a 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.

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 $u_i(x'_i, x_j) > u_i(x_i, x_j)$.

Individual Rationality. A player is called *individually rational at x* in X if she *does not* make a non-profitable deviation from it.

Collective Rationality. A player is called *collectively rational* if she evaluates each possible strategy profile by the minimum payoff she might receive under any individually rational behavior of the other players. Assuming that the others do the same, she aims to maximize that minimum gain by collective maximization principle, that is, by Pareto optimality.

The assumptions of individual and collective rationality are meant for capturing conflicting and cooperating interests of the players respectively. In strategic games, there is always room for both conflict and cooperation unless the game is of zero-sum or the preferences of the players fully coincide. Trying selfishly to maximize utility will be usually of no use because players control only one of the variables leading to the outcomes in games. In general, players do need to cooperate to achieve a certain outcome, but of course this should not be done blindly. Collective rationality captures the cooperating preferences of the players while it respects conflicting preferences among them through individual rationality.

The remaining assumptions are as follows: (a) Each player is individually rational and each player assumes that the others are individually rational. (b) Each player is collectively rational and each player assumes that the others are collectively rational. (c) Each player knows her payoffs and each player assumes that the others know everybody's payoffs. (d) Players do not have any cognitive or computational limitations and each player assumes that the others do not have these limitations. (e) The assumptions (a), (b), (c) and (d) are common knowledge.¹¹

¹¹Curiously enough, the earliest text we found that emphasizes the difference between mutual knowledge and common knowledge (without explicitly using the term) in games is

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 strategic 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 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 $v_1 = \max_{x'_1 \in X_1} v_{x'_1}$. In other words, the unique utility that Alfa can guarantee by facing the uncertainty of playing this game is v_1 . Accordingly, it is recommended that Alfa should choose a strategy $x_1^* \in \arg \max_{x'_1 \in X_1} v_{x'_1}$ which guarantees the value v_1 .

We would like to extend von Neumann's method in such a way that Alfa takes into account the individual rationality of Beta when making the evaluations and vice versa. Let us construct the approach we take step by step and state its implications. We have proposed a notion of individual 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 $v_1(x_1, x_2)$ of the uncertainty for keeping the agreement (x_1, x_2) can be defined as the minimum utility he would receive under any rational behavior of Beta.

the thought-provoking book of Schelling (1960, p. 109, 279, 281). See Lewis (1969) for a detailed discussion and see Aumann (1976) for a formal definition of common knowledge in a Bayesian setting.

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

Let us define the value function formally.

Definition 1. 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 $x = (x_i, x_j) \in X$, the i 'th component of $v = (v_i, v_j)$ satisfies

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

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

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

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

As a consequence, it is not in general true for a strategy $x'_2 \neq x_2$ that we have the equality $v_1(x_1, x_2) = v_1(x_1, x'_2)$. 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 cannot assign a unique value to every strategy of Alfa anymore. Instead, the evaluation of the uncertainty can be encoded in the strategy profile as in the value notion of Nash (1950). Nash defines the value of the game (henceforth the Nash-value) to a player as the payoff that the player receives from a Nash equilibrium when all the Nash equilibria lead to the same payoff for the player. We extend Nash's value approach to the full domain of the game, that is, we assign a value to each single strategy profile including, of course, the Nash equilibria. Notice that when a strategy profile is a Nash equilibrium, the value of a player at this profile is precisely her Nash equilibrium payoff. In particular, if the Nash-value exists for a player then the player's value of every Nash equilibria is the Nash-value of that player. As a result of assigning a value to the profiles rather than the strategies, we can no longer refer to 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 as in a Nash equilibrium. 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 similar inferences as him.

To illustrate what a value function of a game looks like, let us consider the game Γ in Figure 3 which is played by Alfa and Beta. Observe that Γ

		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	-3, -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	-3, -2
	D	3, 1	3, 0	0, 0	3, 0

Figure 3: A game Γ and its value function $v(\Gamma)$.

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 her agreement and may unilaterally deviate to A leaving him a payoff of 1. Accordingly, the value of the profile (C,B) to Alfa is 1 as shown in the bottom table in Figure 3. Now, suppose Alfa offers to make an agreement on (B,B). Beta would not fear a unilateral profitable deviation C of Alfa since she gets 80 in that case. Alfa's payoff does not change 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 .

The second and the last step is to make comparisons between the evaluations of the strategy profiles. Since Alfa and Beta are collectively rational they employ the Pareto optimality principle to maximize the value function. Now, let us formally define the maximin equilibrium.

Definition 2. Let (X_1, X_2, u_1, u_2) be a two-player game and let $v = (v_i, v_j)$ be the value function of the game. A strategy profile $x = (x_i, x_j)$ where $i \neq j$ in a two player game Γ is called *maximin equilibrium* if for every player i and all $x' \in X$, $v_i(x') > v_i(x)$ implies $v_j(x') < v_j(x)$.

Going back to the example in Figure 3, observe that the profile (B,B) is the Pareto dominant profile of the value function of the game Γ , so it is a maximin equilibrium with the value of (90, 80). 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 cannot rule out a potential profitable deviation of Beta to the strategy D. If this happens, Alfa would receive -3 which is less than what he would receive if he did not deviate to C. But Beta is also in the exactly same situation. As a result, it seems that none of them would actually deviate from the agreement (B,B).

Regarding the game presented in Figure 1, notice that $(a, b), (c, a)$ and (b, b) are the maximin equilibria and that these maximin equilibria do not change by the addition of ‘*trick*’ strategies as in Figure 2.

We obtain maximin equilibrium by evaluating each single strategy profile in a game. One of the reasons of the extension of Nash (1950)’s value argument is the following. A Nash equilibrium is solely based on the evaluation of 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. Therefore, it is quite questionable whether the Nash-value represents an evaluation of the strategic uncertainty of the *whole* game. Since a Nash equilibrium completely ignores the outcomes that might occur under any other strategy choices of the players no matter how high their utilities are, this ignorance might lead to a disastrous outcome for both players in strategic games. One can see this clearly in the traveler’s dilemma game which is illustrated in Figure 4 and which was introduced by Basu (1994). If players play the unique Nash equilibrium, then they ignore a large part of the game which is mutually beneficial for both of them, but mutually beneficial trade is perhaps one of the most basic principles in economics. At the end of the day, what a self-interested player cares about is the payoff she receives and not the nice property of being sure that her opponent would not deviate had they agreed to play a profile. 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 an uncertain gamble is very high in every state of the world and the outcome of the sure gamble is very low.

In the traveler’s dilemma, the payoff function of a player i if she plays x_i and her opponent plays x_j is defined as $u_i(x_i, x_j) = \min\{x_i, x_j\} + r \cdot \text{sgn}(x_j - x_i)$ for all x_i, x_j in $X = \{2, 3, \dots, 100\}$ where $r > 0$ determines the magnitude of reward and punishment which is 2 in the original game. Regardless of the magnitude of the reward/punishment, the unique strict Nash equilibrium is $(2, 2)$ which is also the unique outcome of the process of iterated elimination

	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

of strictly dominated strategies.

It is shown by many experiments that the players do not on average choose the Nash equilibrium strategy and that changing the reward/punishment parameter r effects the behavior observed in experiments. Goeree and Holt (2001) found that when the reward is high, 80% of the subjects choose the Nash equilibrium strategy but when the reward is small about the same percent of the subjects choose the highest. This finding is a confirmation of Capra et al. (1999). There, play converged towards the Nash equilibrium over time when the reward was high but converged towards the other extreme when the reward was small. 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 almost¹³ any ‘irrational’ strategy but their Nash equilibrium strategy, then they both get strictly more payoff than they would get by playing the Nash equilibrium. Moreover, 2 is the worst reply in all those cases. In fact, the Nash equilibrium is the only profile which has this property in the game. Therefore, it is difficult to imagine that a self-interested player would ever play or expect her opponent to play the strategy 2 in this game unless she is a ‘victim’ of game theory.

To find the maximin equilibria we first need to compute the value of the traveler’s dilemma. The value function of player i is given by

¹³If one modifies the payoffs of the game such that $u_i(x_i, 3) = 2.1$ and $u_i(x_i, 4) = 2.1$ for all i and all $x_i \in \{4, 5, \dots, 100\}$, then one can even remove ‘almost’ from this sentence.

$$v_i(x_i, x_j) = \begin{cases} x_j - 2, & \text{if } x_i > x_j \text{ for } x_i \in X \\ x_i - 3, & \text{if } x_i = x_j \text{ for } x_i \in X \setminus \{2\} \\ 2, & \text{if } x_i = x_j = 2 \\ x_i - 5, & \text{if } x_i < x_j \text{ for } x_i \in X \setminus \{4, 3, 2\} \\ 0, & \text{if } x_i < x_j \text{ for } x_i \in \{4, 3, 2\}. \end{cases}$$

Observe that the global maximum of the value function is $(97, 97)$ which is assumed at $(100, 100)$. Hence, the profile $(100, 100)$ is the unique maximin equilibrium and $(97, 97)$ is the value of it. It may be interpreted as an ‘ideal’ point for players to reach an agreement (explicitly or implicitly) by respecting the conflicting interests of each other. It is then an empirical question whether players stay within their individually rational behavior with respect to this point. Note that as the reward parameter r increases, the value of the maximin equilibrium decreases. When r is higher than or equal to 50, the unique maximin equilibrium becomes the profile $(2, 2)$ which is also the unique Nash equilibrium of the game. This seems to explain both the convergence of play to $(100, 100)$ when the reward is small and the convergence of play to $(2, 2)$ when the reward is big. Note also that when the reward parameter converges down to 1 the value of the profile $(100, 100)$ converges to 100, and at the limit the profile $(100, 100)$ becomes a Nash equilibrium. It is, however, never a Nash equilibrium for a $r > 1$, although the economic significance of these outcomes seems to barely change when there is a small change in r .

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 1. *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 $B_j(x) = \hat{B}_j(x)$, that is

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

It implies that $\arg \min_{x'_j \in B_j(x)} u_i(x_i, x'_j) = \arg \min_{x'_j \in \hat{B}_j(x)} \hat{u}_i(x_i, x'_j)$ such 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 . 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 maximin equilibrium in pure strategies. This is especially a desired property in games where players cannot 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 1. *Every finite game has a maximin equilibrium in pure strategies.*

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

For another illustrative example, let us consider the game in Figure 5 played by Alfa and Beta. It can be interpreted as the prisoner's dilemma game with a silence option. Each prisoner has three options to choose from, namely 'stay silent' (S), 'deny' (D) or 'confess' (C) and let the payoffs be as in Figure 5. Notice that if the strategy 'stay silent' is removed from the game for both players then we would obtain the prisoner's dilemma whose maximin equilibrium is the same as its Nash equilibrium. Observe that the game has a unique Nash equilibrium (C,C) which is also the unique rationalizable strategy profile (Bernheim, 1984 and Pearce, 1984) with a payoff vector of (10, 10). The maximin equilibria in are (S,S), (D,S) and (S,D) whose values are (100, 100), (5, 110) and (110, 5) respectively.

	Stay silent	Deny	Confess
Stay silent	100, 100	110, 105	0, 15
Deny	105, 110	95, 95	5, 380
Confess	15, 0	380, 5	10, 10

Figure 5: Modified prisoner’s dilemma.

Suppose that the prisoners Alfa and Beta are in the same cell and they can freely discuss what to choose before they submit their strategies. However, they will make their choices in separate cells, that is, non-binding pre-game communication is allowed. A potential agreement in this game seems to be the maximin equilibrium (S,S). By playing her part of the maximin equilibrium, Beta simply guarantees a payoff of 100 under any individually rational behavior of Alfa, and vice versa.

4 The mixed extension of games

4.1 Maximin equilibrium

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 over the set X_i .¹⁴ It is assumed that the preferences of the players over the strategy profiles satisfy weak order, continuity and the independence axioms.¹⁵ As a result, those preferences can be represented by von Neumann-Morgenstern (expected) utility functions $u_1, u_2 : \Delta X_1 \times \Delta X_2 \rightarrow \mathbb{R}$. A mixed strategy profile is denoted by $p \in \Delta X$ 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 Definition 1 and in Definition 2 as being mixed. 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 following proposition shows that maximin equilibrium has this property.

¹⁴For a detailed discussion of the mixed strategy concept, see Luce and Raiffa (1957, p. 74)’s influential book in game theory.

¹⁵For more information see, for example, Fishburn (1970).

Proposition 2. *The maximin equilibria of a game in mixed extension is unique up to positive linear transformations of the payoffs.*

We omit the proof since it follows essentially the same steps as the proof of Proposition 1. The following lemma illustrates a useful property of the value function of a player.

Lemma 1. *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.

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

$Gr(B_j)$ is open in $\Delta X \times \Delta X_j$ if and only if its complement is closed. Let $[(p_j, q_i, q_j)]_{k=1}^{\infty}$ be a sequence in $[Gr(B_j)]^c = (\Delta X \times \Delta X_j) \setminus Gr(B_j)$ converging to (p_j, q_i, q_j) where $p_j^k \notin B_j(q^k)$ for all k . That is, 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 $\hat{u}_i : \Delta X_i \times \Delta X_j \times \Delta X_j \rightarrow \mathbb{R}$ by $\hat{u}_i(q_i, q_j, p_j) = u_i(p_j, q_i)$ for all $(q_i, q_j, p_j) \in \Delta X_i \times \Delta X_j \times \Delta X_j$. 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)¹⁶ 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.¹⁷ 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.

As a result, 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

¹⁶We follow the terminology, especially the definition of upper hemi-continuity, presented in Aliprantis and Border (1994, p. 569).

¹⁷We use the fact that a function f is lower semi-continuous if and only if $-f$ is upper semi-continuous.

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

Theorem 2. *Every finite game in mixed extension has a maximin equilibrium.*

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 1. 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 maximin equilibria in the game. Similarly, one may show that the set v_{ji}^{max} is also a non-empty compact subset of the set of the maximin equilibria. \square

Now, let us assume that players can use mixed strategies in the game Γ in Figure 3. 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 ones. 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 $(90, 80)$.¹⁸ 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 and also that ordinal preferences remain the same. From economics viewpoint the question arises: Should ceteris paribus effect of increasing the payoff of $u_1(C, D)$ from -4 to -3 be substantially high with respect to the solutions of the two games? According to maximin equilibrium the answer is negative. For instance, 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 $\hat{\Gamma}$.¹⁹

4.2 Zero-sum games

Two-player zero-sum games are both historically and theoretically important class in game theory. We illustrate the relationship between the equilibrium

¹⁸The 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 $(88.11, 1.14)$.

¹⁹Note that we have given one example of maximin equilibrium which seems reasonable and whose value is equal for both players, but there can be other maximin equilibria as well.

solution of von Neumann (1928) and the maximin equilibrium in these class of games. The following lemma will be useful for the next proposition.

Lemma 2. *Let (Y_1, Y_2, u_1, u_2) be a two-player zero-sum game with arbitrary strategy sets. Then $v_i(p_i, p_j) = \inf_{p'_j \in Y_j} u_i(p_i, p'_j)$ for each $i \neq j$.*

Proof. Suppose that there exists $\bar{p}_j \in Y_j$ such that $\bar{p}_j \in \arg \min_{p'_j} u_i(p_i, p'_j)$. Then $v_i(p_i, p_j) = \min_{p'_j} u_i(p_i, p'_j) = u_i(p_i, \bar{p}_j)$. Suppose, otherwise, that for all $p'_j \in Y_j$ there exists $p''_j \in Y_j$ such that $u_i(p_i, p''_j) < u_i(p_i, p'_j)$. It implies that $v_i(p_i, p_j) = \inf_{p'_j: u_i(p_i, p'_j) < u_i(p_i, p_j)} u_i(p_i, p'_j) = \inf_{p'_j} u_i(p_i, p'_j)$. \square

The following proposition shows that a possibly mixed strategy profile is a maximin equilibrium if and only if it is a pair of maximin strategies in zero-sum games.

Proposition 3. *Let (Y_1, Y_2, u_1, u_2) be a two-player zero-sum game with arbitrary strategy sets. A profile $(p_1^*, p_2^*) \in Y_1 \times Y_2$ is a maximin equilibrium if and only if $p_1^* \in \arg \max_{p_1} \inf_{p_2} u_1(p_1, p_2)$ and $p_2^* \in \arg \max_{p_2} \inf_{p_1} u_2(p_1, p_2)$.*

Proof. ‘ \Rightarrow ’ Firstly, we show that if (p_1^*, p_2^*) is a maximin equilibrium then its value must be Pareto dominant in a zero-sum game. By contraposition, suppose that (\hat{p}_1, \hat{p}_2) is another maximin equilibrium and suppose without loss of generality that $v_1(p_1^*, p_2^*) > v_1(\hat{p}_1, \hat{p}_2)$ and $v_2(p_1^*, p_2^*) < v_2(\hat{p}_1, \hat{p}_2)$. By Lemma 2, we have $v_1(p_1^*, p_2^*) = v_1(p_1^*, \hat{p}_2)$ and $v_2(\hat{p}_1, \hat{p}_2) = v_2(p_1^*, \hat{p}_2)$. It implies that the value of (p_1^*, \hat{p}_2) Pareto dominates the value of (p_1^*, p_2^*) which is a contradiction to our supposition that (p_1^*, p_2^*) is a maximin equilibrium. Since the value of (p_1^*, p_2^*) is Pareto dominant, each strategy is a maximin strategy of the respective players.

‘ \Leftarrow ’ By Lemma 2 $p_i^* \in \arg \max_{p_i} \inf_{p_j} u_i(p_i, p_j)$ implies $v_i(p_i^*, p_j^*) \geq v_i(p_i, p_j)$ for all $p_i \in Y_i$ and $p_j \in Y_j$. Since the value of (p_1^*, p_2^*) is Pareto dominant it is a maximin equilibrium. \square

As a result, maximin equilibrium indeed generalizes the maximin strategy concept of von Neumann (1928) from zero-sum games to nonzero-sum games. Proposition 3 also shows that maximin equilibria in a deterministic game is not necessarily the same as maximin equilibria in the mixed extension.

Corollary 1. *Maximin equilibrium and equilibrium coincide whenever an equilibrium exists in a zero-sum game.*




Figure 6: The game $(\Delta X, \Delta X_l \cup \Delta X_r, u, -u)$.

For an illustrative example let us consider the following game to be played by Alfa and Beta at a television program. Initially, Beta has to make a choice between the left door and the right door. She is not allowed to commit to a randomization device nor is she allowed to use a device by herself for this choice. If she picks the left door, they will play the game at the left of Figure 6. If she picks the right door, they will play the game at the right of Figure 6. At this stage, players may commit to mixed strategies by submitting them on a computer. Alfa will not be informed which normal-form game he is playing. This situation can be represented by the zero-sum game $(\Delta X, \Delta X_l \cup \Delta X_r, u, -u)$ in which Alfa chooses a mixed strategy in ΔX and Beta chooses a mixed strategy in either ΔX_l or in ΔX_r .

Notice that there is no equilibrium in this game. There are, however, maximin strategies for each player that are $(\frac{11}{12}, \frac{1}{12}) \in \Delta X$ guaranteeing $\frac{-1}{12}$ and $(0, 1) \in \Delta X_l$ guaranteeing 0. By Proposition 3, this pair is also the unique maximin equilibrium whose payoff vector is $(\frac{-1}{12}, \frac{1}{12})$. However, maximin equilibrium does not necessarily say that this is the payoff that players should expect by playing their part of the maximin equilibrium. Rather, the unique maximin equilibrium value of this game is $(\frac{-1}{12}, 0)$. In other words, the unique value of the game to Alfa is $\frac{-1}{12}$ given the individual rationality of Beta and the unique value of the game to Beta is 0 given the individual rationality of Alfa. If the television programmer modifies the game so that Beta is allowed to commit to a randomization device in the beginning, then the game would have an equilibrium $[(\frac{11}{12}, \frac{1}{12}), (0, \frac{11}{12}, 0, \frac{1}{12})]$ which is also a maximin equilibrium. Note that Beta is now able to guarantee the payoff $\frac{1}{12}$. As a result, the unique value of the modified game would be $(\frac{-1}{12}, \frac{1}{12})$.

Speaking of the importance of committing to mixed strategies, let us consider the following zero-sum game on the left in Figure 7 which was discussed

	L	R		L	R
L	0, 0	2, -2		11, 11	13, 6
R	3, -3	1, -1		16, 5	12, 10

Figure 7: Two ordinally equivalent games.

in Aumann and Maschler (1972). Suppose that players cannot commit playing mixed strategies but a randomization device, e.g. a coin is available. Before the coin toss, the maximin strategy $(\frac{1}{2}, \frac{1}{2})$ of Alfa guarantees the highest expected payoff of 1.5. However, after the coin toss Alfa still needs to make a decision whether playing according to the outcome of the toss or not. Actually, for both players playing strategy R guarantees more than playing L after the randomization. Hence the maximin equilibrium of this deterministic game is (R,R) whose value is $(1, -2)$ whereas the values of the profiles (L,L),(L,R) and (R,L) are $(0, -3)$, $(0, -2)$ and $(1, -3)$ respectively. Note, however, that if the utilities of a zero-sum game are measured in an ordinal scale, then the usual intuition of zero-sum games may not hold. For example, the game on the right in Figure 7 is ordinally equivalent to the zero-sum game on the left.

4.3 The relation of maximin equilibrium with the other concepts

Solution of zero-sum games have very strong justifications which can be instructive for nonzero-sum games. The properties of these solutions can be briefly stated as follows: (i) If (p_k, p_l) and (p'_k, p'_l) are two solutions then (p_k, p'_l) and (p'_k, p_l) are also solutions; (ii) The payoff vectors of the all solutions are the same; (iii) If p_k is a solution strategy of player k then it guarantees the solution payoff to k no matter what the other player does; (iv) A player can get at most his solution payoff if the opponent plays her solution strategy; (v) There is no strict incentive for individually rational players to unilaterally deviate from it; (vi) There is no strict incentive for collectively rational players to jointly deviate from the solution.

The property (v) is a unilateral argument for equilibrium and the property (vi) is a collective argument for equilibrium. These two arguments together with the value and the interchangeability arguments make the solution of zero-sum games quite exceptional. It is also perhaps the combination of

these two equilibrium arguments that makes general economic equilibrium so remarkable.

When it comes to nonzero-sum games, it is clearly not possible to incorporate all the six properties in one solution concept. It has long been a subject of discussion which of these properties should a solution of nonzero-sum games possess. Von Neumann and Morgenstern (1944) has given the priority to the value argument for the extension of their solution concept from zero-sum to nonzero-sum games. Nash (1950) has given the priority to the unilateral equilibrium argument and has ingeniously shown that every n -person non-cooperative game has a strategy profile which has this property. He made use of the value and the interchangeability arguments secondarily to distinguish the ‘better’ Nash equilibria from the ‘worse’ Nash equilibria. Nash (1950) calls a game solvable if all the Nash equilibria are interchangeable. Moreover, he defines the upper value of the game to a player as the maximum payoff she gets from a Nash equilibrium and the lower value to a player as the minimum payoff she gets from a Nash equilibrium. Accordingly, the Nash-value of the game to a player is the payoff that she gets from a Nash equilibrium when the upper value equals the lower value.

Consider the ‘best’ situation in which a game has a unique Nash equilibrium so that it is solvable and the game has a value in the sense of Nash (1950) for each player. Given that the Nash-value of the game is entirely based on the evaluation of the outcomes that might occur under one strategy choice of each player, it is not a priori (neither a posteriori) clear whether the Nash-value represents an evaluation of the whole game or only of those outcomes. In contrast to the Nash-value, the maximin strategy value is based on the full evaluation of the uncertainty of the game. Although it is a pessimistic evaluation, we a priori know that this value represents the game as a whole in a pessimistic viewpoint.

In this paper, we extend the value argument and the collective equilibrium argument from zero-sum games to nonzero-sum games. Besides, maximin equilibrium extends the value argument of Nash (1950) to the full domain of the game. When the Nash-value exists for a player, the player’s value of every Nash equilibria is the Nash-value of that player. By evaluating each single strategy profile in a game and by making a comparison between those evaluations we obtain maximin equilibrium. As a result, maximin equilibrium captures both conflicting and cooperating preferences of the players. To put it into other words, a strategy profile is a maximin equilibrium if there is no incentive for collectively rational players to jointly deviate from it.

Let us state Nash (1950)'s path-breaking theorem formally: Every finite game in mixed extension possesses at least one strategy profile p such that $p_i \in \arg \max_{p'_i \in \Delta X_i} u_i(p'_i, p_j)$. The following two propositions illustrate the Pareto dominance relation between Nash equilibrium and maximin equilibrium.

Proposition 4. *For every Nash equilibrium that is not a maximin equilibrium there exists a maximin equilibrium that Pareto dominates it.*

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

Proposition 5. *A maximin equilibrium is never Pareto dominated by a Nash equilibrium.*

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

The two propositions above are closely linked but one does not follow from the other. Because, Proposition 4 does not exclude the existence of a Nash equilibrium that is both Pareto dominated by a maximin equilibrium and Pareto dominates another maximin equilibrium. Proposition 5 shows that this is not the case.

Let us now consider a simple game in Figure 8 to illustrate the difference between maximin equilibrium, rationalizability and correlated equilibrium (Aumann, 1974). This game is nothing else but a strategic decision making situation in which the preferences of the both players can be represented by a single preference relation, i.e. $(l, l) \succ (l, r) \sim (r, r) \sim (r, l)$. Given that both players strictly prefer the outcome (l, l) to the all other outcomes and that the preferences of them over the outcomes are exactly the same, it seems that 'rational' players should do nothing but choose left (l) to achieve this outcome. The unique maximin equilibrium in this game is (l, l) confirming this intuition. It is also a Nash equilibrium, but there is another Nash equilibrium (r, r) which is also a rationalizable strategy profile, a correlated equilibrium and its payoff $(0, 0)$ is a rational expectation according to Aumann and Dreze (2008). The profile (r, r) is usually justified as follows. If

	l	r
l	1, 1	0, 0
r	0, 0	0, 0

Figure 8: A game where $(l, l) \succ (l, r) \sim (r, r) \sim (r, l)$.

Alfa believes that Beta will play r then playing r is best response, and if Beta believes that Alfa will play r then playing r is a best response for Beta and the circular arguments continue ad infinitum. However, the question is: “Do the ‘rational’ players have any reason to believe that the other will play r , given that the preferences of the players are common knowledge?”. The answer of maximin equilibrium is not positive; it predicts that the unique maximin equilibrium value of this game is 1 for each player.

Note that Nash equilibrium, rationalizable strategy profiles and correlated equilibrium are not a generalization of maximin equilibrium since maximin equilibrium is not necessarily an outcome of the iterated elimination of strictly dominated strategies. As discussed earlier, the profile (2,2) is the only outcome of this process in the traveler’s dilemma, but it is not a maximin equilibrium.

One might wonder whether there is a relationship between the maximin (minimax) decision rule²⁰ in decision theory and the maximin equilibrium. Imagine a one-player 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 decision maker 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.

5 Maximin equilibrium in n -person games

The analysis of n -person games significantly differ from those of two-person games. This is partially due to the possibility that players may form coalitions.

²⁰See Wald (1950) for maximin decision rule and see Gilboa and Schmeidler (1989) for an axiomatization of it.

	D	E
A	1, 1, 1	0, 0, 1
B	4, 6, 2	0, 4, 6
C	-2, 1, 1	0, 0, 2

	D	E
A	-2, -1, 6	3, 3, -2
B	3, 4, 3	-1, 5, 8
C	4, 0, 0	-1, 0, 1

Figure 9: A three player game where player 3 chooses between the matrices L (left) and R (right).

tions. We firstly present the n -person extension of the definition of maximin equilibrium in ignorance of coalitional deviations.

Regarding the definition of the value function, we need to replace the way v_i is written in Definition 2 to $v_i(p) = \min\{\inf_{p'_{-i} \in B_{-i}(p)} u_i(p_i, p'_{-i}), u_i(p)\}$ where $B_{-i}(p)$ is defined as follows. Firstly for each $S \subseteq N \setminus \{i\}$ and each $p \in \Delta X$ define

$$B_{-i}^S(p) = \{(\hat{p}_S, p_{-S}) \in \Delta X_{-i} \mid u_k(\hat{p}_k, p_{-k}) > u_k(p) \text{ for all } k \in S\}.$$

$B_{-i}^S(p)$ is the set of $(n - 1)$ -tuple strategy profiles in which the players in S make a *unilateral* profitable deviation with respect to p . To represent all such profiles for all $S \subseteq N \setminus \{i\}$, we define the correspondence $B_{-i}(p) = \bigcup_{S \subseteq N \setminus \{i\}} B_{-i}^S(p)$. Accordingly, a strategy profile is a maximin equilibrium if its value is not Pareto dominated. Moreover, every result in Section 3 and in Section 4 except Proposition 3 (which requires a zero-sum setting) is valid in n -player games. The proofs are essentially the same as the ones given in Section 3 and in Section 4.

Let us consider the game in Figure 9 played by Alfa, Beta and Juliet in order to show that the unique Nash equilibrium may break down in games with more than two players. This game has a unique Nash equilibrium which is approximately $[(0.65, 0.35, 0), (0.25, 0.75), (0.68, 0.32)]$ whose payoff vector is approximately $(0.71, 2.12, 2.39)$. Note that the Nash-value of Juliet is the highest so she seems to be the most advantageous player in the game. Suppose that Juliet naively thinks that she is doing the ‘best’ by playing her part of the Nash equilibrium. Even without any communication, Alfa and Beta may unilaterally deviate from the Nash equilibrium to the strategies B and C respectively and can gain strictly more than their Nash equilibrium payoff which causes the Nash equilibrium to collapse. Moreover, Juliet ends up with a strictly less payoff than her payoff at the Nash equilibrium. Notice also that Alfa and Beta make the unilateral deviations at the cost of nothing, because

B of Alfa is a best response to the Nash equilibrium strategies of the other players and D of Beta is also best response to the Nash equilibrium strategies of the others. Besides, holding the Nash equilibrium strategy (0.68, 0.32) of Juliet fixed, the profile (B,D) is a Nash equilibrium in the game played by Alfa and Beta! Consequently, the very argument that players have no incentive to unilaterally deviate at a Nash equilibrium breaks down because Alfa and Beta simply do have incentives to unilaterally deviate. Since every pure strategy in the support of a mixed Nash equilibrium is a best response, every mixed Nash equilibrium and even sometimes a pure Nash equilibrium may, potentially, have the problem described above.²¹

In fact, von Neumann and Morgenstern (1944, p. 32) anticipate this situation years before the emergence of Nash equilibrium: “Imagine that we have discovered a set of rules for all participants to be termed as “optimal” or “rational” each of which is indeed optimal provided that the other participants conform. Then the question remains as to what will happen if some of the participants do not conform. If that should turn out to be advantageous for them and, quite particularly, disadvantageous to the conformists then the above “solution” would seem very questionable. We are in no position to give a positive discussion of these things as yet but we want to make it clear that under such conditions the “solution,” or at least its motivation, must be considered as imperfect and incomplete.”

Maximin equilibrium can be modified to incorporate coalitions in n -person games, we just need to define the the better reply correspondence allowing coalitional profitable deviations and define the value function with respect to this. Accordingly, a profile is called strong maximin equilibrium if its value is not Pareto dominated. By the same argument in Theorem 1, it exists in pure strategies in the deterministic game. Regarding the mixed extension of games, one may show the existence of strong maximin equilibrium by following the similar steps as in Lemma 1 and in Theorem 2. Regarding the three-player game above, both the maximin equilibrium and the strong maximin equilibrium is the profile (B, D, $(\frac{1}{2}, \frac{1}{2})$) whose value is (3, 4, 2.5). In

²¹It is well-known that a Nash equilibrium may break down because some group of players (coalition) may make a profitable deviation from it. Therefore some refinements of Nash equilibrium has been proposed such as strong Nash equilibrium (Aumann, 1959) and coalition-proof Nash equilibrium (Bernheim et al., 1987). These concepts, however, have the non-existence problem and they are sometimes interpreted with pre-play communication. In the game presented in Figure 9, we illustrate a simple game in which the unique Nash equilibrium breaks down and it happens even without any communication.

other words, by playing their part of the maximin equilibrium each player guarantees her value under any unilateral or coalitional profitable deviation of the other players.

6 Conclusion

In this paper, we extended von Neumann's theory of games from zero-sum games to nonzero-sum games by incorporating common knowledge of individual and collective rationality of the players. The assumptions of individual and collective rationality are meant for capturing the conflicting and the cooperating interests of the players, respectively. Maximin equilibrium extends Nash's value approach to the whole game and evaluates the strategic uncertainty of the game by following a similar method as von Neumann's maximin strategy notion. We show that maximin equilibrium is invariant under strictly increasing transformations of the payoffs. Notably, every finite game possesses a maximin equilibrium in pure strategies. We showed that maximin equilibrium is invariant under strictly increasing transformations of the payoff functions of the players. Notably, every finite game possesses a maximin equilibrium in pure strategies.

Considering the games in von Neumann-Morgenstern mixed extension, we demonstrated that maximin equilibrium exists in mixed strategies as well. Moreover, we showed that a strategy profile is a maximin equilibrium if and only if it is a pair of maximin strategies in two-person zero-sum games. In particular, a maximin equilibrium and an equilibrium coincide whenever they exist in those games. Furthermore, we show that for every Nash equilibrium that is not a maximin equilibrium there exists a maximin equilibrium that Pareto dominates it. Besides, a maximin equilibrium is never 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 such as the further exploration in extensive form games and in repeated games.

References

Aliprantis, C. and K. Border (1994). *Infinite Dimensional Analysis: A Hitchhiker's Guide*.

- Aumann, R. J. (1959). Acceptable points in general cooperative n -person games. In R. D. Luce and A. W. Tucker (Eds.), *Contribution to the theory of game IV, Annals of Mathematical Study 40*, pp. 287–324. University Press.
- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics 1*(1), 67–96.
- Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics 4*(6), pp. 1236–1239.
- Aumann, R. J. (1990). Nash equilibria are not self-enforcing. *Economic Decision Making: Games, Econometrics and Optimisation*, 201–206.
- Aumann, R. J. and J. H. Dreze (2008). Rational expectations in games. *American Economic Review 98*(1), 72–86.
- Aumann, R. J. and M. Maschler (1972). Some thoughts on the minimax principle. *Management Science 18*(5-Part-2), 54–63.
- Basu, K. (1994). The traveler’s dilemma: Paradoxes of rationality in game theory. *The American Economic Review 84*(2), 391–395.
- Berge, C. (1959). *Espaces topologiques: Fonctions multivoques*. Dunod.
- Bernheim, B., B. Peleg, and M. D. Whinston (1987). Coalition-proof Nash equilibria i. concepts. *Journal of Economic Theory 42*(1), 1 – 12.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica 52*(4), pp. 1007–1028.
- Capra C. M., Holt C. A., J. K. G. and R. Gomez (1999). Anomalous behavior in a traveler’s dilemma? *American Economic Review 89*(3), 678–690.
- Fishburn, P. (1970). *Utility theory for decision making*. Publications in operations research. Wiley.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.
- Goeree, J. K. and C. A. Holt (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review 91*(5), 1402–1422.

- Harsanyi, J. and R. Selten (1988). *A General Theory of Equilibrium Selection in Games*. Mit Press.
- Harsanyi, J. C. (1966). A general theory of rational behavior in game situations. *Econometrica* 34(3), pp. 613–634.
- Lewis, D. (1969). *Convention: a philosophical study*. Harvard University Press.
- Luce, R. and H. Raiffa (1957). *Games and Decisions: Introduction and Critical Survey*. Dover books on advanced mathematics. Dover Publications.
- McKelvey, Richard D., M. A. M. and T. L. Turocy (2013). Gambit: Software tools for game theory, version 14.0.1.
- Nash, J. F. (1950). Non-cooperative games. PhD Diss. Princeton University.
- Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics* 54(2), 286–295.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4), pp. 1029–1050.
- Rubinstein, A. (2006). Dilemmas of an economic theorist. *Econometrica* 74(4), 865–883.
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal* 117(523), 1243–1259.
- Savage, L. (1954). *The Foundations of Statistics*. Wiley publications in statistics. John Wiley and Sons.
- Schelling, T. (1980). *The Strategy of Conflict*. Harvard University Press.
- von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* 100, 295–320.
- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior* (1953, Third ed.). Princeton University Press.
- Wald, A. (1950). *Statistical decision functions*. Wiley publications in statistics. Wiley.
- Wolfram Research, I. (2010). Mathematica version 8.0.