TI 2004-028/1
Tinbergen Institute Discussion Paper

A Discussion of Maximin

Vitaly Pruzhansky

Faculty of Economics and Business Administration, Vrije Universiteit Amsterdam, and Tinbergen Institute.

Tinbergen Institute
The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Please send questions and/or remarks of non-scientific nature to driessen@tinbergen.nl.
Most TI discussion papers can be downloaded at http://www.tinbergen.nl.

A DISCUSSION OF MAXIMIN

VITALY PRUZHANSKY

Abstract. This paper builds on one of the results of Pruzhansky [22], namely that maximin strategies guarantee the same expected payoffs as mixed Nash equilibrium strategies in bimatrix games. We present a discussion on the applicability of maximin strategies in such class of games. The usefulness of maximin is illustrated from both positive and normative viewpoints. Examples are provided.

1. Introduction

It is not a coincidence that the title of this paper parallels that of Owen [21]. Likewise, we provide a discussion on the applicability of maximin strategies in completely mixed bimatrix games, and discuss their (potential) superiority as compared to mixed equilibrium strategies. This work continues the theme started in our recent paper Pruzhansky [22], to which the reader is referred for precise notation and definitions, such as maximin and equalizer strategies.

Recall how a typical justification of Nash equilibria runs: if somebody is to recommend equilibrium strategies to the players, then no one has incentives *not* to obey this recommendation. When it comes to completely mixed equilibrium strategies, this justification is somewhat weakened: even though no player has incentives to deviate, they still have no particular incentives to randomize with the prescribed probabilities. Nevertheless, despite the fact that such equilibrium strategies typically do not guarantee expected equilibrium values to the players, they rule out the possibility of strategic outguessing by the opponent, and thus may seem desirable. However, a more reasonable description of reality would probably involve players figuring out themselves what to do in a particular game independently of each other. Suppose a player discovers that he has a maximin strategy that guarantees him the same expected payoff as the equilibrium one, although it leaves open

Date: February 20, 2004.

Key words and phrases. Bounded rationality, common knowledge of rationality, correlated equilibria, rationalizability, uncertainty aversion.

I am thankful to J. Kamphorst and G. van der Laan.

the possibility that the opponent will correctly predict the player's intentions and respond accordingly[1]. What strategy should such a player choose? In other words, in what cases can the desire to guarantee the expected equilibrium payoff be at least as important as ruling out possible strategic outguessing[2]? Alternatively, if we were to estimate the chances that the outcome of a particular game will be consistent with players' selecting maximin or mixed equilibrium strategies, what strategy would have higher odds?

The purpose of this paper is to answer the above questions and to support the claim that maximin strategies may be at least as attractive as the mixed equilibrium ones, and, thus, should be taken into account[3]. This claim will be defended from both normative and positive standpoints. Section 2 presents the positive arguments while Section focuses on the normative ones. Finally, Section 4 provides examples illustrating each of these positions.

It is essential to note that most of our arguments are based on the idea that players are not certain about the strategies chosen by their opponents. Consequently, they treat these strategies as random variables, to which subjective probability distributions must be assigned. This approach is not new and was used in Kadane and Larkey [16] to question the power of the minimax solution in two person zero-sum games. Their position received a fair amount of criticism in Harsanyi [14] due to the fact that it may not agree with the basic principles of game theory as spelled out by J. von Neumann and O. Mongerstern. Nevertheless, we believe that in *non*-zero sum games the force of Harsanyi's critique is limited. In these games players' interests are not always directly opposed. That is, player i's objective is not necessarily to minimize the expected payoff of player j, but to maximize his own. Consequently, maximization of own expected utility may conflict with the desire to avoid being outguessed by the opponent. This phenomenon is typical in games, whose Nash equilibria are only mixed.

Before we start the discussion, let us restate what is understood by the term 'mixed strategy', since numerous interpretations of this terminology exist. For all Sections the interpretation of mixed strategy as a *belief*, in the spirit of Aumann [2] would suffice. However, in Sections

[1] To sharpen our arguments, in what follows we suppose that the set of strategies that are both equilibrium and maximin is empty.

[2] This should not be confused with obtaining an equilibrium payoff *for sure*, since a maximin strategy may be mixed!

[3] We are not aware of any applied model, in which the distinction between mixed equilibrium and maximin strategies is clearly made. Typically, maximin strategies are simply disregarded.

3.1 - 3.2 it is also useful to take another view, a classical one, according to which mixed strategies are considered to be the object of choice. It will be fair to say that this approach has often had conceptual difficulties and probably other approaches, such as the above interpretation in terms of beliefs or that of Harsanyi [12], are more intuitive. Nevertheless, as we show by a few examples in Section 3.1 below, sometimes players may have strict incentives to randomize in order to get the maximum of their expected utilities.

2. Positive Arguments

2.1. **Common Knowledge.** The use of maximin strategies in the class of games we consider is consistent with the assumption of common knowledge of rationality. This follows from the fact that in completely mixed bimatrix games no player has dominated strategies. Thus, all pure and mixed strategies are rationalizable in the sense of Bernheim [5], and maximin strategies are rationalizable too. Unfortunately, rationalizability of maximin strategies cannot be generalized to arbitrary (bimatrix) games. If the game under consideration is *not* completely mixed, then a maximin strategy of one player may be rationalizable only via a dominated strategy of the opponent.

From the point of view of common knowledge of Bayesian rationality (see Aumann [2] for precise definition) when all players have subjective priors, we are led to the concept of *subjective* correlated equilibrium. Maximin strategies are the part thereof, since rationalizability is equivalent to subjective correlated equilibria in the class of completely mixed games, as follows from the work of Branderburger and Dekel [6]. One has to remember, though, that it is not the desire to maximize the minimum of expected utility that 'forces' players to select their maximin strategies in a subjective correlated equilibrium. It is the difference in priors and information partitions. Also notice that the critique of rationalizability and subjective correlated equilibria voiced in Aumann [2], p. 14, does not apply to completely mixed games.

However, if the assumption of common knowledge of Bayesian rationality is accompanied by the requirement that players' priors be the same, it follows from Aumann [2] that some *objective* correlated equilibrium must take place. The latter no longer uniformly supports maximin strategies. For instance, in 2×2 games with a unique Nash equilibrium in mixed strategies the only correlated equilibrium is identical to the completely mixed one (see Calvó-Armengoly [7]). For games of higher dimensions, to the best of our knowledge, no results have been obtained yet.

Thus, the common prior assumption turns out to be crucial for justification of maximin strategies, at least for 2×2 games. A discussion on whether this assumption is plausible can be found in Aumann [2], pp. 12-15. Note, however, that even if maximin strategies correspond to some (objective) correlated equilibria, Nash mixed strategies always do too. Hence, although it cannot be claimed that maximin strategies are *always* better than mixed equilibrium ones, at least rationality does *not* violate the use of maximin.

2.2. **Bayesian Beliefs in 2×2 Games.** Bimatrix games, in which each player has only two pure strategies are, perhaps, the most well-known and widely applied class of games. We will show how Bayesian rationality supports players' use of their maximin strategies in completely mixed games of this specific dimension. For application of Bayesian rationality in a more general setting see Tan and Werlang [23]. One difference of our analysis and that of [23] is that we focus specifically on mixed equilibrium and maximin strategies. Interestingly, the framework we develop here turns out to have an unexpected relation with the notion of risk dominance.

Let x be an n-dimensional random variable, whose realizations are restricted to Δ, and F (f) its multivariate probability distribution (joint density function). Thus,

$$\Pr[x \in S \subseteq \Delta] = \int_S f(s)\, ds,$$

where ds stands for $ds_1 ... ds_n$ and $\int_S$ for the n-dimensional integral over the region S.

Define the expected value of x as the vector of expectations, i.e.

$$E[x] = (E[x_1], ..., E[x_n]),$$

such that for all $i = 1, ..., n$

$$E[x_i] = \int_\Delta s_i f(s)\, ds.$$

(In what follows, the region of integration will always be confined to the unit simplex, thus, the symbol Δ will be suppressed). As a shorthand for this formula we will write

$$E[x] = \int s dF.$$

Suppose that the decision choice of player 2 is perceived by player 1 as a random n-dimensional vector[4] q, and the beliefs of player 1 are described by a probability distribution F over Δ. Let $\mathcal{F}$ be a set of possible beliefs of player 1 and suppose that all members of $\mathcal{F}$ have distinct means (see footnote 5 for an explanation). That is, for any $x \in \Delta$ there exists a unique $F \in \mathcal{F}$ satisfying

$$x = \int s dF.$$

Clearly, there can be many sets $\mathcal{F}$ satisfying the above condition. However, any such set will serve our purposes.

Define a metric $d(\cdot)$ on $\mathcal{F}$ in the following way. Let $x = \int sdF$ and $y = \int tdG$ for some $F, G \in \mathcal{F}$, then

$$d(F,G) = \left[\sum_{i=1}^{n} (x_i - y_i)^2\right]^{\frac{1}{2}}.$$

That is, the distance between two distributions F and G is given by the usual Euclidian distance between their mean values. It is easy to check that the above metric $d(\cdot)$ satisfies all required axioms of metric[5].

The expected payoff of player 1 under belief F when he plays strategy p is

$$v_1 = \int (pAs)\, dF = pA\left(\int sdF\right). \tag{2.1}$$

In principle, F may be degenerate in the sense that its density f puts all probability mass on a single point $\widetilde{s}$. This would correspond to player 1 believing that player 2 chooses $\widetilde{s} \in \Delta$ with certainty. Under such belief F (2.1) reduces to the usual $v_1 = pA\widetilde{s}$.

For $p^* \in \Delta$, define $\mathcal{F}^* \subset \mathcal{F}$ as follows

$$\mathcal{F}^* := \left\{F \in \mathcal{F} |\ p^* \in \arg\max_p \left(pA \int sdF\right)\right\}.$$

Using the properties of a completely mixed equilibrium (p^*, q^*), it is easy to verify that for all $F \in \mathcal{F}^*$ it must be the case that $q^* =$

[4]Strictly speaking, it must be an $n-1$ dimensional vector, because variables $q_1, ..., q_n$ are linearly dependent. However, to simplify notation we will talk about n dimensions.

[5]The non-negativity property $d(x,y) = 0$ if and only if $x = y$ forced us to assume that all members of $\mathcal{F}$ have distinct mean values. This can be interpreted as if decision-makers cared only about expected values, and not about higher moments of random payoffs. Moreover, they regard two different distributions F and G with the same mean as being identically equal.

$\int sdF$. Therefore, $\mathcal{F}^*$ represents the set of beliefs, to which p^* is a best response.

Similarly, define $\overline{\mathcal{F}} \subset \mathcal{F}$

$$\overline{\mathcal{F}} := \left\{ F \in \mathcal{F} | \ \overline{p} \in \arg\max_p \left(pA \int tdF \right) \right\}$$

to be the set of player 1's beliefs justifying the use of $\overline{p}$.

Finally, let $\mathcal{F}_i \subset \mathcal{F}$ be

$$\mathcal{F}_i = \left\{ F \in \mathcal{F} | \ e_i \in \arg\max_p \left(pA \int udF \right) \right\}.$$

Thus, $\mathcal{F}_i$ is the set of beliefs under which the optimal response of player 1 is a pure strategy $p = e_i$.

Since in completely mixed games there are no dominated strategies, any (pure or mixed) strategy of every player is a best response to some beliefs. Hence, the sets $\mathcal{F}^*$, $\overline{\mathcal{F}}$, $\mathcal{F}_i$ for each $i = 1, ..., n$ are non-empty. Similarly, any mixed strategy p of player 1 is justifiable, given some belief F (and this is trivially so for $F = F^*$). Suppose that p assigns positive probabilities to some $k \leq n$ pure strategies of player 1, without loss of generality let them be the first k strategies. Due to the linearity of expected payoffs in bimatrix games, it follows that any pure strategy in this set $\{1, ..., k\}$ must be also justifiable by F. Thus,

$$F \in \mathcal{F}_1 \cap ... \cap \mathcal{F}_k,$$

and this intersection in non-empty. We summarize these relations in a Lemma below, which, because of its simplicity, is stated without proof.

Lemma 1. *Let (A, B) be a completely mixed (possibly non-generic) bimatrix game and let the maximin strategy of player $\overline{p}$ assign positive probabilities only to his $k \leq n$ pure strategies. Then it holds that*

$$\begin{aligned} \overline{\mathcal{F}} &\equiv \bigcap\nolimits_{i=1}^{k} \mathcal{F}_i, \\ \mathcal{F}^* &\equiv \bigcap\nolimits_{i=1}^{n} \mathcal{F}_i, \\ \mathcal{F}^* &\subseteq \overline{\mathcal{F}}. \end{aligned}$$

Moreover, in generic games $\mathcal{F}^* \equiv \{F^*\}$, such that $q^* = \int sdF^*$. Furthermore, if $\overline{p}$ is completely mixed, then $\mathcal{F}^* \equiv \overline{\mathcal{F}} \equiv \{F^*\}$.

Next Lemma will be used in the proof of the subsequent theorem.

Lemma 2. *Let (A, B) be a completely mixed generic bimatrix game. For any set of $k < n$ pure strategies of player 1, there exists a belief $F \neq F^*$, such that any strategy $p \in \Delta$, assigning positive probability only to the above k pure strategies is a best response, given F.*

Proof. Fix a not completely mixed strategy $\widetilde{p}$ of player 1. We need to show the existence of $\widetilde{q}$ such that

$$\widetilde{p}A\widetilde{q} \geq pA\widetilde{q}, \tag{2.2}$$

for all $p \in \Delta$.

Consider the following system

$$\widetilde{A}q = \widetilde{u}, \tag{2.3}$$

where

$$\widetilde{A} = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{k1} & \cdots & a_{kn} \\ \gamma a_{k+11} & \cdots & \gamma a_{k+1n} \\ \vdots & & \vdots \\ \gamma a_{n1} & \cdots & \gamma a_{nn} \end{pmatrix}, \ \gamma > 1,$$

and all coordinates of $\widetilde{u}$ are equal. Since the game is generic, $\det(A) \neq 0$ and it is easy to check that $\det\left(\widetilde{A}\right) \neq 0$ too. Thus, there exists a unique solution to (2.3), say $\widetilde{q}$. Now observe that given F, such that $\widetilde{q} = \int s dF$, the strategy $\widetilde{p}$ assigning positive probabilities only to the first k rows of A solves (2.2). We just need to ensure that $\widetilde{q} \in \Delta$. Indeed,

$$\lim_{\gamma \to 1} \widetilde{q} = q^*.$$

Hence, for γ close enough to 1, we will have that $\widetilde{q} \in \Delta$. □

The intuition for Lemma 2 is simple. The existence of a completely mixed equilibrium implies the existence of a unique hyperplane passing via points $a_1, ..., a_n$ with a normal q^*. If the vectors $a_1, ..., a_n$ are linearly independent, we can always 'lift' the hyperplane so that it rests on only k points $a_1, ..., a_k$ and lies above the other $n-k$ points. The 'shifted' hyperplane, thus, has a normal $\widetilde{q}$ that solves (2.2). An important consequence of this Lemma is that any *not* completely mixed strategy $\widetilde{p}$ must be a *strict* best response to $q = \alpha\widetilde{q} + (1-\alpha) q^*$ for any $\alpha \in (0, 1]$, i.e.

$$\widetilde{p}Aq > pAq, \text{ for all } p \neq \widetilde{p}.$$

That is, for any non-completely mixed $\widetilde{p}$ there exists a continuum of distributions, whose mean values lie on a line segment between $\widetilde{q}$ and q^*, to which $\widetilde{p}$ is a best response. Formally, for each not completely mixed strategy $\widetilde{p}$ there exists $\widetilde{F} \in \mathcal{F}$, such that for any $G \in \mathcal{F}$ enjoying

$$d\left(\widetilde{F}, G\right) \leq \epsilon,$$

where ϵ is sufficiently small, it holds that

$$\widetilde{p}A\int sdG \geq pA\int tdG, \text{ for all } p \in \Delta.$$

This implies the following.

Theorem 1. *Let $\mathcal{A}$ be any σ-algebra of $\mathcal{F}$, containing all $\mathcal{F}_i$, $i = 1, ..., n$, and $\mathcal{G}$ be the class of all generic completely mixed bimatrix games. For any probability measure $\mu(\cdot)$ on $\mathcal{A}$ and any game $(A, B) \in \mathcal{G}$, it holds that $\mu(\mathcal{F}^*) \leq \mu\left(\overline{\mathcal{F}}\right)$. Moreover, if the measure $\mu(\cdot)$ is non-atomic*[6] *and the maximin strategy $\overline{p}$ of player 1 is not completely mixed, then $0 = \mu(\mathcal{F}^*) < \mu\left(\overline{\mathcal{F}}\right)$.*

Proof. By Lemma 8, $\mathcal{F}^* \subset \overline{\mathcal{F}}$, thus for any probability measure μ we have $\mu(\mathcal{F}^*) \leq \mu\left(\overline{\mathcal{F}}\right)$, with equality if and only if $\overline{p}$ is completely mixed. Because of the genericity assumption, $\mathcal{F}^* \equiv \{F^*\}$. On the other hand, if $\overline{p}$ is not completely mixed, then $\overline{\mathcal{F}}$ is not singleton. Hence, for any non-atomic μ we have $\mu(\mathcal{F}^*) = 0$ and $\mu\left(\overline{\mathcal{F}}\right) > 0$. □

The Theorem asserts that player 1 is at least as likely to hold beliefs justifying his use of maximin strategy, as his use of the mixed equilibrium strategy. That is, Bayesian rationality in bimatrix games favors maximin strategies at least as often as it does mixed equilibrium strategies. Moreover, given any $F \in \mathcal{F}^*$, any pure or mixed strategy of player 1 yields him the same expected payoff. Thus, if players are Bayesian expected utility maximizers and are not completely rational in the sense that they disregard the consequences of their own actions on the actions of the opponents, mixed equilibrium strategies will be observed only by chance, perhaps just in the case, when they easily suggest themselves, like in games á la Matching Pennies.

This interpretation of the Theorem is close to the positive support for maximin strategies in a subjective correlated equilibrium. There, if the outside observer witnessed the frequencies with which players choose their strategies, those frequencies would form a subjective correlated equilibrium. Here, the claim is that the outside observer is more likely to witness an outcome induced by players' maximin strategies. Note also, that the formation of players' beliefs is *not* based on past plays, but *is* arbitrary. It is as if players have never met before. Then their intention to play maximin strategies arises much more often than the intention to play a completely mixed Nash equilibrium strategy.

---

[6] Measure μ is said to be non-atomic if for any $A \in \mathcal{A}$ such that $\mu(A) > 0$, there exists $B \subset A$, so that $\mu(B) > 0$. In other words, any finite set of points (in our case distributions) receives zero weight. This assumption reflects the diversity of opinions that player 1 may hold about player 2's possible behavior.

From the above Theorem it also follows that for any pure strategy, to which $\overline{p}$ assigns positive probability, say the k-th one, we have

$$\mu(\mathcal{F}^*) \leq \mu\left(\overline{\mathcal{F}}\right) \leq \mu(\mathcal{F}_k).$$

Thus, Bayesian rationality would, perhaps, favor the k-th pure strategy even more often than $\overline{p}$. Why is the maximin strategy special? Indeed, in general bimatrix games we do not find much support for maximin strategies as compared to other pure strategies. However, in 2×2 completely mixed games the attractiveness of maximin comes in a new light. The reasons is that if $\overline{p}$ is not completely mixed in generic games, then it must be pure. Therefore,

$$0 = \mu(\mathcal{F}^*) < \mu\left(\overline{\mathcal{F}}\right).$$

In words, according to a Bayesian view, player 1 *is* more likely to play a pure maximin strategy than a completely mixed equilibrium one for any set of beliefs $\mathcal{F}$, if μ is non-atomic. This is an interesting fact, taking into account the range of applications in which the theory of 2×2 games has been used. There is a need to see if such a hypothesis is confirmed empirically[7].

Finally, the approach of comparing the sets of beliefs $\mathcal{F}_i$ and $\mathcal{F}_j$ turns out to be intrinsically linked to risk dominance in 2×2 games (regardless of whether maximin strategies are pure or not). Consider an arbitrary 2×2 bimatrix game with two distinct *pure* strategy Nash equilibria. For concreteness, assume that both players selecting either their first or second pure strategies constitutes a Nash equilibrium. Such a game will also have a completely mixed Nash equilibrium. Let $\widetilde{q}$ be an equilibrium probability of player 2 selecting his first pure strategy in this mixed Nash equilibrium. Clearly for player 1 we have

$$\begin{aligned}
\mathcal{F}_1 &\equiv \left\{F \in \mathcal{F} | \int sdF \in [\widetilde{q}, 1]\right\}, \\
\mathcal{F}_2 &\equiv \left\{F \in \mathcal{F} | \int tdF \in [0, \widetilde{q}]\right\}.
\end{aligned}$$

Consider now a *uniform* measure[8] μ, i.e. the measure that is proportional to the length of the line segments $\mathcal{F}_1$ and $\mathcal{F}_2$. Suppose that $\widetilde{q} \leq \frac{1}{2}$ and, correspondingly, $\mu(\mathcal{F}_1) \geq \mu(\mathcal{F}_2)$. If a similar relation holds for

[7] In order to confirm it, one probably needs to employ the interpretation of mixed strategies as frequencies, with which corresponding pure strategies are selected over time.

[8] That is, an outside oberver will witness all possible beliefs of player 1 in the set $\mathcal{F}$ equally likely.

player 2, then the equilibrium in which both players select their first pure strategies is risk dominant.

3. Normative Arguments

3.1. **Two Polar Versions of Rationality.** Rationality can take up many forms. At the one extreme players are rational in the Bayesian sense. They form beliefs (which can be arbitrary) and act rationally, given these beliefs. In addition, each player considers the world as fixed at any specific point in time. Maximin strategies of these players are the natural response to the uncertainty surrounding them. At the other end of the spectrum, there lies a case of *complete* rationality, which is characterized not only by common knowledge of rationality, but also by the fact that players fully account how their own actions affect the actions of others. Specifically, a completely rational player i will not take those actions, to which his opponents have incentives to respond in an unfavorable to i manner, as compared to some status quo state. Although we do not define this notion formally, it is not hard to see that a solution concept, appropriate in the presence of complete rationality, should not be just 'individually' rational, but also 'commonly' rational, i.e. it must constitute a Nash equilibrium. In terms of set conclusion, complete rationality implies common knowledge of rationality, and the latter implies Bayesian rationality. Moreover, in completely mixed games Bayesian rationality and common knowledge of rationality coincide[9].

One of the most popular critique of mixed strategy Nash equilibria is that players have to randomize with correct probabilities in order to make the *opponent* indifferent. Statements like 'rational players do not have incentives to randomize with equilibrium probabilities, since they are indifferent' are abundant, eg. Harsanyi and Selten [15], pp. 14-16, Osborne and Rubinstein [20], p. 41. Such claims can be in line with the notion of Bayesian rationality, but are inconsistent with the idea of complete rationality. The reason being that the latter concept is not limited to payoffs only, but also takes into account how actions/intentions of one player affect actions/intentions of his opponents. That is, ruling out strategic outguessing is necessarily a consequence of complete rationality. (In the same time, it may not be implied by the common knowledge assumption, see below). We will argue that completely rational players *do* have incentives to adhere to their mixed

[9]This is due to the absence of dominated strategies in completely mixed games.

equilibrium strategies[10]. However, the conditions that ensure this are very strict and, perhaps, unrealistic. That is why maximin strategies may seem to be better.

Consider the following two games.

	L	R
T	$1, 1$	$3, 0$
B	$2, 0$	$2, 1$

Game Γ_1

	L	M	R
T	$0, 5$	$1, 0$	$2, 2$
B	$1, 0$	$0, 3$	$2, 2$

Game Γ_2

It can be checked that in both games the unique equilibrium strategy of player 1 is completely mixed and is not equal to his maximin strategy. In both cases, thus, player 1 is indifferent and presumably does not have incentives to randomize as the equilibrium prescribes. Moreover, in Γ_1 he can guarantee himself the equilibrium value by selecting his pure strategy B. If player 1 indeed does not care, what player 2 thinks about him, the desire to obtain the equilibrium value with certainty seems to be a rational course of action. However, if player 1 takes into account that player 2 will try to outguess him, he should *not* select B with certainty. Thus, complete rationality forces player 1 to randomize.

In Γ_2 the situation is even more serious. The unique Nash equilibrium of the game prescribes that player 1 randomizes $(p, 1-p)$, where $p \in \left[\frac{1}{3}, \frac{2}{5}\right]$, and player 2 selects R. Although this equilibrium is *not* completely mixed, given player 2's strategy R, player 1 is indifferent between any two pure or mixed strategies. Now, for any $p \notin \left[\frac{1}{3}, \frac{2}{5}\right]$ player 2 will *never* play R. As a consequence, the expected payoff for player 1 will not exceed one under any circumstances! However, in equilibrium, he obtains the payoff of two. Therefore, player 1 has *strict* incentives to randomize. Moreover, his randomization cannot be arbitrary! To push this idea further, observe that Γ_2 can be modified so that the set of admissible equilibrium strategies of player 1 shrinks to

[10]This statement should be treated cautiously. First, what players really must do is to credibly persuade the opponent that they in fact randomize. Second, there might be peculiar exceptions, where the equilibrium value can be obtained with certainty, not employing any randomization, even if the opponents are aware of the deviation, see Examples below.

a single point. Also note that although common knowledge of rationality does not preclude player 1 from selecting $p \notin \left[\frac{1}{3}, \frac{2}{5}\right]$, complete rationality does so! In particular the maximin strategy $\overline{p} = \left(\frac{1}{2}, \frac{1}{2}\right)$ is rationalizable, but if player 1 is known to use it, then player 2 will certainly not play R, and as a consequence, player 1's expected payoff will not exceed one. (Of course, player 1 can obtain his equilibrium payoff of two if player 2 decides to use his own maximin strategy, i.e. R. However, the point we wish to emphasize is that player 1 has strict incentives to randomize, when he is seemingly indifferent).

It may seem that these examples are artificial and in games like the one presented in the Introduction of [22] or in Aumann and Maschler [4] such phenomena do not arise. We, however, are convinced that these differences are not justified. The game from the Introduction is simply an extreme version of the above reasoning. But the reasoning goes along exactly the same lines as in the above examples.

What are the conditions that one needs in order to ensure players' complete rationality and how strict are they? The following condition is sufficient, but it appears quite severe to us[11]. Namely, it must be common knowledge among players that any deviations from equilibrium strategies will be correctly exploited by the opponents. In sum, the adherence to (completely) mixed equilibrium strategy of player i that depends on j's payoffs rests on a very demanding assumption. Whenever the latter is *not* satisfied, mixed Nash equilibria seem to be especially vulnerable in the presence of maximin strategies that can guarantee the same expected payoff.

3.2. **Implementation of Mixed Strategies.** As was demonstrated above, completely rational players may have strict incentives to randomize. In their classical interpretation, mixed strategies are viewed as an object of choice. That is, players are assumed to use a random device, such as a roulette spin, in order to decide about their choice of a pure strategy. If such position is adopted, the problem of *implementation* arises. Recall that an equilibrium mixed strategy of a player is optimal only if the opponent also follows his equilibrium strategy. Therefore, before actually implementing such a strategy one should be convinced that the opponent randomizes according to the equilibrium probabilities himself. That is, he constructs a random experiment

[11]It does not rule out a possibility that some less strict conditions may lead to the same outcome. We have not explored this possibility yet.

and commits himself to the subsequent outcome[12]. Moreover, the randomizing probabilities of the random device must be objective in the sense that all players agree on them. These arguments become especially important when the equilibrium mixed strategy is not obvious (eg. Matching Pennies), but involves a randomization over a big set of strategies that have to be played with different probabilities. In fact, unless a player supervises the experiment conducted by the opponent and observes its outcome, he cannot be convinced entirely that the opponent is playing according to the Nash equilibrium. And, of course, after the outcome was observed, he will probably never conduct an experiment himself, as the best response can be computed right away.

It is noteworthy that maximin strategies do not suffer from similar drawbacks. The decision to use them depends only on the player's desire, and correct implementation is subject only to player's skills. As to the possible criticism that player j will respond optimally to i's maximin if he knows about i's intentions to use it - it is ill grounded. If the above reasoning is valid, namely, players do perceive the issue of implementation to be relevant, then player j simply *cannot* be sure that i will play his maximin strategy correctly, even if the latter is known to have such an intention. In this situation the best player j can do is to resort to his own maximin, especially as the latter guarantees the equilibrium payoff. Related to the problem of implementation and bounded rationality is the criticism of the statement that if player i questions j's ability to compute or implement the equilibrium strategy correctly, i's equilibrium strategy may still be a good choice, since j's mistakes may be 'in i's favor'. It is not hard to see that, given the mistakes of j, the equilibrium strategy of player i can make him either at most as good, as any other strategy can, or strictly worse off.

The problem of subjectivity in players' assessments about each other's randomizing devices was also raised in Aumann [1], that naturally led to the development of subjective correlated equilibrium discussed above. However, Aumann does not himself make any reference to the importance of maximin strategies, which appears very logical to us in this regard[13].

[12]The issue of commitment is often overlooked in textbooks, although it is a very serious point. For an intriguing example of how security considerations *after* randomization may force a player to violate the commitment see Owen [21]. Similar ideas are also stated in Luce and Raiffa [18], pp. 75-76.

[13]Aumann supposes that each player has a subjective probability distribution about the state of the world. However, when players' uncertainty is so big that they cannot assign *any* probability distribution, the use of maximin strategies seems to be the only option.

3.3. **Bounded Rationality.** Empirical evidence largely supports the fact that neither do players trust their opponents to be rational, nor are they capable of computing non-trivial Nash equilibrium strategies, see, for instance, Fudenberg and Tirole [9], p. 8, Goeree and Holt [10][14]. (Note, that even if players can compute their equilibrium strategies, there still remain problems of implementation). How should a mildly rational player, who has the necessary skills to analyze the game correctly, behave in such an interaction then? If he cannot observe empirical frequencies of the opponent's choices and does not have any record of his past behavior, it seems plausible that guaranteeing the equilibrium value should outweigh the desire to outguess the opponent.

Perhaps in the same spirit, although the term 'bounded rationality' was not explicitly used, Harsanyi [13], p. 116, formulates a set of rationality postulates. One of them asserts the following: if player i cannot hope to obtain more than his maximin payoff $\overline{v}_i$, he must use the strategy that 'fully assures at least that much'[15]. Hence, maximin and not mixed equilibrium strategies are defended. It must be stressed that the example, in which Harsanyi illustrates this idea is a *non*-zero sum, completely mixed bimatrix game of the type we deal with in this paper.

3.4. **Multiple Equilibria and Uncertainty Aversion.** Typically, the use of maximin strategies, either pure or mixed, is supported in the framework of subjective or nonadditive probability that was axiomatized by Gilboa and Schmeidler [11], see for instance Lo [17] and Marinacci [19]. These models favor maximin strategies because players understand that they face uncertainty on the part of their opponents' strategies. Depending on the degree of this uncertainty non-maximin strategies become either more or less attractive. As for the maximin ones - they are attractive for any level of uncertainty, since the individuals are *uncertainty averse*. In our setup, a bimatrix game (A, B) may possess multiple mixed strategy equilibria, which naturally leads to the above uncertainty.

[14]Despite the fact that the latter tested common knowledge of rationality only in extensive games of perfect information, their findings can probably be of some interest here as well. Specifically, Goeree and Holt [10], found that players did *not* trust in each other's rationality when the costs of such *irrationality* was small.

[15]It is not hard to see that in any game player j can always hold (perhaps unconsciously) i's expected payoff down to $\overline{v}_i$ by an appropriate randomization.

4. Further Examples

Here we provide two Examples illustrating the ideas of the previous Sections. Roughly speaking, the first Example focuses on positive aspects, while the second stresses normative usefulness of maximin strategy.

4.1. Bounded Rationality and Model Misspecification.

In applied research we rarely know players' exact utility functions. Usually the researcher simply substitutes monetary payoffs for utilities. Nash equilibria of these misspecified models confirm with reality only by chance, of course. Below we present an example, where such model misspecification highlights the use of maximin strategies. It also shows that the strategic situation, as perceived by players, may be different from the model designed to capture it.

We also maintain the assumption that players are boundedly rational. What we specifically mean by this is that players have a certain model in mind and they act in accordance with its solution. These models may differ. It appears that this is much more realistic than simply assuming that all players have the *same* model in mind.

Example 1. *A taxpayer (player 1) decides whether to report his income honestly or evade some part of it. Independently of the taxpayer, a tax agency (player 2) decides whether to audit him or not. Auditing imposes fixed costs on the agency. If the audit is carried out then the true income of the taxpayer is detected with probability one. If cheating indeed took place, the taxpayer pays the missing amount of tax plus a fine that is proportional to the amount of concealed income. The monetary payoffs are summarized in the following table*[16]

	audit	*don't*
evade	$Y(1-t)-\theta(Y-y),\ Yt+\theta(Y-y)-c$	$Y-yt,\ yt$
don't	$Y(1-t),\ Yt-c$	$Y(1-t),\ Yt$

[16] We are not claiming that this overly simple model best describes important aspects of tax evasion. There are much more sophisticated models designed for that. The usefulness of this example lies in the fact that, despite being general enough, it exemplifies two polar cases: when maximin strategies coincide with equilibrium ones, and when they are pure.

Where

Y - *real income of the taxpayer*
y - *reported income*
t - *tax rate*
θ - *penalty rate (per unit of concealed income)*
c - *fixed costs of audit*

It is further assumed that the agency is risk neutral and the taxpayer is risk averse, and the costs of audit are not too high, namely

$$c \leq (Y - y)(t + \theta).$$

The misspecification of this model, therefore, lies only in the wrong representation of the taxpayer's preferences. Simple computations show that

$$\begin{aligned}
\overline{p} &= (0, 1), \\
p^* &= \left(\frac{c}{(Y - y)(t + \theta)}, 1 - \frac{c}{(Y - y)(t + \theta)} \right), \\
q^* &= \overline{q} = \left(\frac{t}{t + \theta}, \frac{\theta}{t + \theta} \right), \\
v_1^* &= Y(1 - t), \\
v_2^* &= Yt - \frac{ct}{t + \theta}.
\end{aligned}$$

There are several points of interest here. First, observe that maximin and equilibrium strategies for the agency coincide[17]. By all means, it would be difficult to recommend any other strategy for the agency in this case. We would like to relate this observation to the comments of Ariel Rubinstein in [20], p. 37, who criticized modeling the relationship between taxpayers and a tax agency as a mixed strategy Nash equilibrium, since in equilibrium the agency is indifferent between any two (not necessarily equilibrium) strategies, which does not seem to be realistic. The fact that $q^* = \overline{q}$ shows that such critique is not entirely convincing: although in equilibrium any strategy yields the same payoff to the agency, only the equilibrium strategy guarantees this payoff independently of the actions of the taxpayer. Hence, the agency may have strong incentives to adhere to q^*.

Second, note that the above v_2^* will be the agency's equilibrium value for *any* strictly increasing concave utility function $u(\cdot)$ applied to the taxpayer's monetary payoffs (as long as the structure of Nash equilibrium is still preserved under this transformation). Therefore, even

[17]Of course, this model is misspecified, but even for a correct model it may turn out that $q^* = \overline{q}$.

despite the fact that the researcher does not know $u(\cdot)$, he may have strong incentives to recommend q^* to the agency, as it guarantees v_2^* independently of the taxpayer's true utility function. The revenue maximizing agency, thus, may adopt the prescribed strategy even if it *suspects* a misspecification, but does not know *how* to amend it. Alternatively, the agency may follow the strategy q^* because of the arguments in Harsanyi [12]. Namely, it may think that the real payoffs of the taxpayer are subject to small perturbations that tend to zero in the limit. And that the population of taxpayers consists of two types, each playing a corresponding pure strategy. This small amount of incomplete information about the type of taxpayer is enough to persuade the agency to randomize according to q^*.

Third, and most important, the model predicts a certain level of taxpayer's evasion, given the strategy of the agency. However, if the agency indeed implements q^*, the risk averse taxpayer has incentives to deviate from the proposed equilibrium. More importantly, he will deviate precisely to $\overline{p}$. To check this, observe that given q^*, the pure strategy *evade* offers the taxpayer a lottery with the expected value $Y(1-y)$, whereas the pure strategy *don't* offers the same amount with certainty. Hence, the real level of evasion will be effectively zero. Furthermore, such deviations to pure maximin strategies would arise even if the latter would *not* guarantee the expected equilibrium payoff. This follows from the fact that pure maximin strategies minimize payoff dispersion (see Lemma 6 in Pruzhansky [22]). It seems that the above misspecification may (at least) partially explain why the levels of taxpayer's evasion in reality are much lower than it is predicted by the theory. Factors summarized in Sections 2 and 3 can also add to this explanation, of course. Moreover, for the reasons explained there, maximin strategy may be preferred to the mixed equilibrium one, even if the model were correctly specified, or the taxpayer were risk neutral.

Finally, on pure psychological grounds, it may be the case that individuals lexicographically prefer a lottery offering v^* with certainty to a lottery, whose expected value is v^*, but which allows receiving strictly less than v^* with some positive probability. We lack precise empirical evidence on that[18]. Similarly, in the context of bounded rationality and one-shot games the quest for certainty may outweigh strategic considerations. These arguments, if confirmed experimentally, will offer yet another support for maximin strategies.

[18] In the context of one-person decision making this is supported by a thought experiment in Ellsberg [8].

4.2. **Unknown Distribution.** Although the next Example does not belong to the class of 2-person games in normal form, it illustrates how considering a maximin strategy may be useful from the normative point of view in the light of uncertainty aversion.

Example 2. *Two players jointly own a business, for which they have private valuations* $\theta_i \in [0,1]$, $i = 1,2$. *Each player* i *possesses a share* s_i *of the total enterprise, thus,* $s_1 + s_2 = 1$. *The utility of player* i *is given by* θ_i, *if he fully owns the business, and zero otherwise. The players decide to split the ownership according to the following procedure. Player 1 names a price* p *for the whole business. Then player 2 either sells his share for* $s_2 p$ *or buys the share of player 1 for* $s_1 p$, *whatever gives him higher utility. What is the equilibrium price?*

In order to solve this problem player 1 needs a probability distribution of θ_2. Let $F(\cdot)$ be such distribution. It is easy to verify that player 1 chooses such p, that maximizes his expected utility, given by

$$v_1(s_1, p) = \Pr[2 \text{ buys}]\, s_1 p + \Pr[2 \text{ sells}]\,(\theta_1 - (1 - s_1)\, p).$$

Clearly, player 2 buys whenever $\theta_2 \geq p$ and sells otherwise. Thus,

$$v_1(s_1, p) = s_1 p - pF(p) + \theta_1 F(p). \tag{4.1}$$

This expression can be maximized with respect to p, after $F(p)$ is assumed a specific functional form. For instance, if $F(p)$ is uniform on $[0,1]$, we have

$$p^* = \frac{s_1 + \theta_1}{2},\ v_1^* = \frac{(s_1 + \theta_1)^2}{4}.$$

Certainly this problem cannot be solved without a specific assumption about $F(\cdot)$. In reality player 1 will never know $F(\cdot)$ for sure. Moreover, let us assume that player 1 is so uncertainty averse that he is afraid of making any choice between possible distributions. At first glance the problem does not have a solution then. Nevertheless, we will see that analyzing a maximin strategy may (though not always!) significantly simplify the choice of player 1.

A maximin solution makes player 1 indifferent between buying or selling of player 2. Thus, he should set $\overline{p} = \theta_1$, which will give him utility $\overline{v}_1 = s_1 \theta_1$ with certainty. It can be checked that $v_1^* \geq \overline{v}_1$, and equality holds only if $s_1 = \theta_1$. Player 1 can use the maximin pricing rule when he is uncertainty averse. However, we can say something more as well. Let us see when this maximin solution also maximizes (4.1). By differentiating (4.1) with respect to p and setting $p = \theta_1$, one finds the following condition

$$s_1 = F(\theta_1), \tag{4.2}$$

which is nothing less than $\Pr[\theta_2 \leq \theta_1] = s_1$, using the definition of $F(\cdot)$. Thus, if player 1 can estimate the probability that his valuation exceeds that of player 2 as being equal to s_1, then he can safely use the maximin pricing and guarantee himself the equilibrium level of utility. It appears that answering this question is somewhat simpler than assigning a *right* probability distribution to θ_2. Of course, it may not solve the problem completely. In particular, generically for an arbitrary distribution F, condition (4.2) will *not* hold with equality. Despite this, considering the maximin price may be a useful step in analyzing a problem like the above.

References

[1] Aumann, R. (1974) "Subjectivity and Correlation in Randomized Strategies", *Journal of Mathematical Economics*, 1, 67-96.

[2] Aumann, R. (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality", *Econometrica,* 55, 1-18.

[3] Aumann, R., A. Brandenburger (1995) "Epistemic Conditions for Nash Equilibrium", *Econometrica*, 63, 1161-80.

[4] Aumann, R., M. Maschler (1972) "Some Thoughts on the Minimax Principle", *Management Science,* 18, 54-63.

[5] Bernheim, B. (1984) "Rationalizable Strategic behavior", *Econometrica,* 52, 1007-27.

[6] Brandenburger, A. E. Dekel (1987) "Rationalizability and Correlated Equilibria", *Econometrica*, 55, 1391-1402.

[7] Calvó-Armengoly, A. (2003) "The Set of Correlated Equilibria of 2×2 Games", available at *selene.uab.es/acalvo/correlated.pdf.*

[8] Ellsberg , D. (1961) "Risk, Ambiguity and the Savage Axioms", *Quarterly Journal of Economics*, 75, 643-69.

[9] Fudenberg, D., J. Tirole *Game Theory*, MIT Press 1991.

[10] Goeree, J., C. Holt (2001) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions", *American Economic Review*, 91, 1402-22.

[11] Gilboa, L., D. Schmeidler (1989) "Maximin Expected Utility with Non-unique Prior", *Journal of Mathematical Economics*, 18, 141-53.

[12] Harsanyi, J. (1973) "Games with Randomly Distributed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points", *International Journal of Game Theory,* 2, 1-23.

[13] Harsanyi, J. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations,* Cambridge Univ. Press 1977.

[14] Harsanyi, J. (1982) "Comments on Kadane and Larkey's Paper", *Management Science,* 28, 120-24.

[15] Harsanyi, J., R. Selten *A General Theory of Equilibrium Selection in Games,* MIT Press 1988.

[16] Kadane, J., P. Larkey (1982) "Subjective Probability and the Theory of Games", *Management Science,* 28, 113-20.

[17] Lo, K. C. (1996) "Equilibrium in Beliefs Under Uncertainty", *Journal of Economic Theory*, 71, 443-84.

[18] Luce, R. D., H. Raiffa *Games and Decisions, Introduction and Critical Survey*, John Wiley&Sons, Inc. 1957.

[19] Marinacci, M. (2000) "Ambiguous Games", *Games and Economic Behavior*, 31, 191-219.

[20] Osborne, M., A. Rubinstein *A Course in Game Theory*, MIT Press 1994.

[21] Owen, G. (1974) "A Discussion of Minimax", *Management Science,* 20, 1316-17.

[22] Pruzhansky, V. (2003) "Maximin Play in Two-person Bimatrix Games", *Tinbergen Institute Discussion Paper No.* 2003-101/1.

[23] Tan, T., S. Werlang (1988) "The Bayesian Foundations of Solution Concepts of Games", *Journal of Economic Theory*, 45, 370-91.

FEWEB, Vrije Universiteit, De Boelelaan 1105, 1081 HV, Amsterdam, The Netherldans

E-mail address: pruzhansky@tinbergen.nl