Rationalizability under Uncertainty using Imprecise Probabilities[*]

Hailin Liu
Department of Philosophy
Carnegie Mellon University, PIttsburgh
hailinl@andrew.cmu.edu

Abstract

The notion of imprecise probability can be viewed as a generalization of the traditional notion of probability. Several theories and models of imprecise probability have been suggested in the literature as more appropriate representations of uncertainty in the context of single-agent decision making. In this paper I investigate the question of how such models can be incorporated into the traditional game-theoretic framework. In the spirit of *rationalizability*, I present two new solution concept called Γ-*maximin rationalizability* and E-*rationalizability*. They are intended to capture the idea that each player models the other players as decision makers who all employ Γ-maximin or E-admissibility as their decision rules. Some properties of these concept such as existence conditions and the relationships with rationalizability are studied.

Keywords. Normal form games, imprecise probabilities, rationalizability, Γ-maximin, E-admissibility.

1 Introduction

The theory of subjective expected utility (axiomatized by Savage [1954]) has become a widely-accepted normative theory for dealing with single-agent decision making under uncertainty. However, the assumption about the representation of uncertainty in this framework has often been criticized for being overly restrictive. In particular, Ellsberg [1961] has argued that uncertainty, as opposed to risk, cannot be adequately represented by a single personal probability distribution. Inspired by this challenge, various alternative theories of decision making under uncertainty have been developed in the literature, e.g., Gilboa and Schmeidler's multiple priors model [1989] and Schmeidler's Choquet expected utility model [1989]. In addition, there has been a vast amount of literature on alternative approaches to representing uncertainty in decision problems, such as upper and lower probabilities, sets of probability measures, belief functions, and so on. (See [Walley, 1991] for a detailed discussion of the models of imprecise probabilities)

The Ellsberg paradox arises in single-agent decision making situations where uncertainty regarding some exogenous event is involved. Nevertheless, one would expect that similar situations of uncertainty could arise in multi-agent, interactive scenarios, where the considerations underlying uncertainty for each player are the other players' strategy choices, rather than the state of nature. This naturally suggests a new line of research, which is to incorporate some model of uncertainty using imprecise probabilities into traditional game-theoretic frameworks. New conceptual issues arise in this approach to game under uncertainty, e.g., how should solution concepts be defined given the new decision theoretic foundations. In recent years, there has been a growing literature on applying the aforementioned theories of imprecise probabilities in the context of games, which can be divided roughly into two categories depending on the way of addressing these conceptual issues. On the one hand there are those that investigate the consequences of allowing players' beliefs to be represented by imprecise probabilities in the framework of Nash equilibrium or its refinements. Dow and Werlang [1994] introduce an equilibrium concept for two-player normal form games in which players' beliefs about the opponents' strategy choices are represented by non-additive probabilities and players are Choquet expected utility maximizers. Eichberger and Kelsey [2000] extend Dow and Werlang's equilibrium concept to normal form games with n-players and discuss some nice properties of this concept. By using the multiple priors model to represent players' uncertainty, Klibanoff [1994] and Lo [1996] provide two equilibrium-type solution concepts

---

for normal form games with any finite number of players. Unlike these researchers, Liu [2011] presents a different solution concept called *robust equilibrium* by extending the framework of the so-called *linear tracing procedure* (Harsanyi and Selten [1988]), to accommodate games with uncertainty where players' initial beliefs are modeled by a set of probability measures rather than a common prior. This concept can be viewed as a refinement of Nash equilibrium.

On the other hand, several studies have attempted to generalize the concept of *rationalizability* (Bernheim [1984] and Pearce [1984]) in normal form games to accommodate notions of rationality other than subjective expected utility maximization. In addition to the idea of equilibrium with uncertainty aversion, another significant innovation introduced by Klibanoff [1994] is the characterization of common knowledge of rationality under uncertainty for normal form games where each player attempts to maximize the minimum expected utility. Epstein [1997] also considers normal form games and develops a general framework for discussing the implications of common knowledge of rationality in which the definition of rationality can accommodate different kinds of preference structures including the multiple priors model.

The approach to game theory with uncertainty I present in this paper is very much in the same spirit as Klibanoff's and Epstein's approaches, which embraces the essential idea of rationalizability, namely, to assume that each player models the opponents as the same kind of rational decision maker under uncertainty. As noted in previous literature, rationalizability captures the idea that each player attempts to deduce their opponents' rational behavior from the structure of the game by modeling her opponents as expected utility maximizers, where players' uncertainty about their opponents' strategy choices are fully described by a single probability distribution. This paper explores the possibility of adapting this standard assumption by using a set of probability distributions to model uncertainty in normal form games. However, even in single-agent decision theory, there is no generally accepted criterion for decision making under uncertainty when uncertainty is depicted by a set of probability distributions. In view of this, this paper develops a general theoretical framework to analyze the implications of rationality and common knowledge of rationality in the sense that each player employs the same decision rule to choose the best strategy with respect to a set of probability distributions. In particular, I consider here two familiar decision rules named Γ-*maximin* [Berger, 1985; Gilboa and Schmeidler, 1989] and *E-admissibility* [Levi, 1974]. According to the former

rule, a decision maker should choose an option that maximizes the minimum expected utility with respect to a set of probability distributions, while the latter one constrains the decision maker's admissible choices to those options that maximizes expected utility for some probability in the set of probabilities. In analogy with rationalizability, I put forward two game-theoretic solution concepts under uncertainty, in which each player is required to model the other players as the same kind of decision makers who use either Γ-maximin or E-admissibility to make decisions. This gives rise to the two solution concepts that we shall call Γ-*maximin rationalizability* and *E-rationalizability* respectively. Just as Γ-maximin and E-admissibility are extensions of subjective expected utility theory, both Γ-maximin rationalizability and E-rationalizability turn out to be generalizations of rationalizability. Example 1 in Section 4 illustrates the distinction between these three solution concepts.

The main contribution of this paper is in providing a general game-theoretic framework which enables us to discuss how different decision rules can be incorporated into the framework of rationalizabiity in normal form games when uncertainty is depicted by a nontrivial set of probability distributions. This framework can be easily adapted to accommodate some other decision rules discussed in decision theory such as Maximality [Walley, 1991]. Although it turns out that the concept of Γ-maximin rationalizability coincides with Klibanoff's and Epstein's iterative definitions of rationalizability with uncertainty aversion (they used different terms for this concept), the current approach to rationalizability under uncertainty can be regarded as complementary work to their theories, since it provides an alternative way of characterizing the same solution concept. By applying this new definition, it is easier to check whether a strategy of a player is Γ-maximin rationalizable (or uncertainty aversion rationalizable). In a similar way, I define the concept of E-rationalizability, which, to my knowledge, has not been explored in any previous study.

The rest of this paper proceeds as follows: Section 2 presents a brief review of the solution concept rationalizability, and discusses some of its properties. Section 3 motivates the idea of using imprecise probabilities to represent uncertainty in games. I then propose two solution concepts called Γ-maximin rationalizability and E-rationalizability, which extend the framework of rationalizability to contexts where a set of probabilities is used to represent uncertainty. Section 4 studies some properties of these solution concepts, and also includes an example to illustrate their difference. Section 5 concludes the paper and suggests possible future work.

2 Rationalizability

In contrast with the concept of Nash equilibrium where each player's belief is required to coincide with her opponents' strategies, the concept of rationalizability, proposed independently by Bernheim [1984] and Pearce [1984], imposes a weaker requirement on players' beliefs. More precisely, it only demands players to obey the requirement of Bayesian rationality and common beliefs in Bayesian rationality. It attempts to account for rational behavior as the consequence of common knowledge of the game structure and the rationality of players, without imposing any further constraints on players' strategy choices.

Let us begin with some formal notations and definitions. Throughout this paper, we consider a finite normal or strategic form game $G \equiv \langle I, \{S_i\}, \{u_i\} \rangle_{i \in I}$, where $I = \{1, 2, \ldots, n\}$ is a finite set of players, S_i denotes a finite set of pure strategies (or actions) available to player i, and $u_i : S \to \mathbb{R}$ denotes player i's payoff function. We shall denote the set of player i's mixed strategies by Δ_i, which can be regarded as the set of all probability distribution over S_i. For each mixed strategy $\delta_i \in \Delta_i$, let $\delta_i(s_i)$ denote the probability assigned to s_i. Recall that a strategy profile is a *Nash equilibrium* if no player can benefit by merely changing her strategy while the other players keep theirs unchanged. More precisely, a mixed strategy profile $\delta^* \in \Delta$ is a (mixed strategy) Nash equilibrium if for each player i, $u_i(\delta_i^*, \delta_{-i}^*) \geqslant u_i(\delta_i, \delta_{-i}^*)$ for every mixed strategy δ_i of player i. An alternative way to characterize the notion of Nash equilibrium is to define it in term of best response. We say that a strategy $\delta_i \in \Delta_i$ is a *best response* to δ_{-i} for player i if $u_i(\delta_i, \delta_{-i}) \geqslant u_i(\delta_i', \delta_{-i})$ for all $\delta_i' \in \Delta_i$. Thus a strategy profile is a Nash equilibrium if each player's strategy is a best response to the other players' strategies. For an arbitrary set X of strategies, we denote by $\mathcal{H}(X)$ the convex hull of the set X, namely, the smallest closed convex set containing X.

It is well known that the concept of rationalizability attempts to characterize rational strategic behavior that are consistent with the assumption that both the structure of the game and the rationality of the players are common knowledge to them. To be more specific, rationalizability in normal form games is defined based on the following assumptions:

- **A1**: Each player employs a subjective personal probability to express her belief about the other players' strategy choice, which cannot conflict with any information available to her.

- **A2**: Each player attempts to maximize expected utility with respect to her subjective probability

regarding her opponents' strategy choices.

- **A3**: The structure of the game, including the strategy space and payoff functions, and the fact that each player satisfies the above two assumptions are common knowledge.

Informally speaking, we can examine a player's rationality by checking whether the actions chosen by that player are "rational" or not. We say that an action of a player is rational if there exists some belief regulated by the assumptions given above such that it is a best response to that belief. Thus, a strategy δ_i of player i is *rationalizable* if she can justify her choice by explaining that (i) δ_i is rational, (ii) there exists some belief δ_{-i} such that δ_i maximizes her own expected utility with respect to δ_{-i}, and δ_{-i} assigns positive probability only to rational actions of her opponents, and (iii) there are beliefs of her opponents that make those actions rational and assign positive probability only to her rational actions, and so on. This suggests an intuitive way of defining rationalizability without invoking the iterative process originally suggested by Pearce (1984). In order to present this formal definition, we have to make the notion of a belief and what we mean by a strategy being rational explicitly.

Definition 1. *In a strategic form game G, a* belief *of player $i \in I$, denoted by μ_{-i}, about the other players' strategy choices is a probability distribution over the set of the other players' strategies $S_{-i} \equiv \prod_{j \neq i} S_j$.*

Here we should draw a clear distinction between the concepts of belief and mixed strategy. A belief about player i has the same mathematical form as a mixed strategy of player i, which is normally found in the literature. However, the interpretations of both concepts are different (see [Osborne and Rubinstein, 1994] for a comprehensive discussion on the interpretations of mixed strategies.). A mixed strategy of player i is usually viewed as an explicit randomization over her pure strategies in S_i. If player i chooses to play a mixed strategy, she commits herself to carry out the deliberate randomization. The main criticism of this interpretation of mixed strategy is that for each player there are usually infinitely many mixed strategies that yield her the same expected payoff as her mixed strategy equilibrium does, given her opponents' equilibrium behavior. But we are here concerned with a different solution concept called rationalizability. Thus, interpreting mixed strategies as objects of deliberate choice is appropriate within the current framework. On the other hand, a belief about player i is a probability distribution on the set of player i's mixed strategies, which represents another player's view about player i's strategy choice. It should not be confused with a randomization that is actually carried

out by player i. In that sense, we can say that players' mixed strategies should be understood as the objects of the beliefs about players' strategy choices, and the probability distribution given by a belief about player i merely represents the likelihood that another player i assigns to player i's mixed strategies.

Nevertheless, an essential feature of this formulation of belief is that it allows a player to believe that the other players choose their strategies according to certain correlated randomization devices, since a belief μ_{-i} of player i is a probability measure over S_{-i} and thus is an element of the set $\mathcal{H}(S_{-i})$. Note that a belief μ_{-i} of player i is not necessarily a product of independent probability distributions on each of the set S_j of actions for $j \in N \backslash \{i\}$. That is, a belief μ_{-i} of player i need not be identified as an element of the set of mixed strategies of her opponents S_{-i}. In addition, it is not difficult to see that the set S_{-i} is strictly smaller than the set $\mathcal{H}(S_{-i})$ in games with more than 2 players. Hence we have to use a different notation μ_{-i} for a belief in the current framework in order to distinguish it from a mixed strategy δ_{-i}.

It is assumed that each player always chooses an action to maximize her own expected payoff with respect to some belief about the opponents' strategies. A strategy being rational can then be defined precisely in terms of maximization of expected utility.

Definition 2. *A strategy δ_i of player i in a strategic form game G is a* rational *strategy if there exists a belief μ_{-i} of player i such that δ_i maximizes player i's expected utility, that is, $u_i(\delta_i, \mu_{-i}) \geqslant u_i(\delta_i', \mu_{-i})$ for all $\delta_i' \in \Delta_i$. In this case, we say that δ_i is a* best response *to the belief μ_{-i}.*

The key idea of the following characterization is to define an action (or pure strategy) to be rationalizable by considering each player's introspective process of justifying her own strategy choice, based on the analysis of her opponents' similar reasoning about their rational behavior.

Definition 3. *In a strategic form game G, an action $s_i \in S_i$ of player i is rationalizable if for each player $j \in I$, there exists a set $Z_j \subseteq S_j$ of actions such that: (i) $s_i \in Z_i$, and (ii) every action s_j in Z_j is a best response to some belief μ_{-j} of player j whose support is a subset of Z_{-j}.*

Whenever a new solution concept is put forward, a primary theoretical question is whether the proposed concept can give rise to at least one solution for games in general. Regarding the concept of rationalizability, the answer to this question is positive.

Proposition 2.1 (Pearce, 1984)**.** *For finite normal form games, the set of rationalizable strategies is al-*

ways nonempty and contains at least one pure strategy for each player.

We have considered above how to define the concept of rationalizability by using the notion of belief and the rationality of the players. As a matter of fact, the set of rationalizable actions can be further characterized for finite strategic games in terms of the familiar idea of dominance relations. As we shall see, this characterization for rationalizability gives rise to an operationalizable method for finding the set of rationalizable actions for finite games. Recall that the concept of rationalizability basically captures the idea that as a rational decision maker each player can only choose those strategies that are best responses to some beliefs regarding the other players' strategies. In other words, a rational player should not adopt a strategy that is not a best response to any belief about her opponents' strategy choices. In the game-theoretic terminology, such a strategy is called a never-best response strategy. Thus one can see that the concept of rationalizability is closely related to the notion of never-best response strategy as defined below.

Definition 4. *In a normal form game G, an action s_i of player i is a* never-best response *if it is not a best response to any belief of player i, that is, for every belief μ_{-i} of player i there exists a strategy $\delta_i \in \Delta_i$ such that $u_i(\delta_i, \mu_{-i}) > u_i(s_i, \mu_{-i})$.*

In other words, there is no belief μ_{-i} of player i about her opponents' strategies with respect to which a never-best response action s_i maximizes her own expected payoff. This coincides exactly with the central idea of rationalizability, namely that the players are rational in the sense of maximizing expected utility. As mentioned above, each player should rule out the actions that are not best response to any belief, namely, never-best response actions.

Let us now turn to the familiar notion of strict dominance which will play a crucial role in the characterization of rationalizable actions, as we shall see below.

Definition 5. *In a normal form game G, an action s_i of player i is* strictly dominated *if there exists a strategy $\delta_i \in \Delta_i$ such that $u_i(\delta_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.*

In words, whatever the other players do, player i can benefit from playing some other strategy rather than a strictly dominated strategy. Clearly, a rational player would never use a strictly dominated strategy. Otherwise the player's choice violates the assumption of rationality in the sense of maximizing expected utility. At this point one may wonder whether the notion of never-best response is equivalent to the conception of strictly dominated action. It turns out that one can

establish the equivalence between these two notions within the current framework.

Lemma 2.2 (Pearce, 1984). *In a strategic form game G, an action s_i^* of player i is a never-best response if and only if s_i^* is strictly dominated.*

Suggested by the above lemma, we can show that the set of rationalizable actions can be obtained by iteratively deleting strictly dominated actions until we arrive at the stage where no more strictly dominated action can be further eliminated. Let us first formally define the process of iterated elimination of strictly dominated actions.

Definition 6. *Consider a normal form game G. Set $S_i^0 \equiv S_i$ for each $i \in I$. Then, for each $i \in I$ and for each $k \geqslant 1$, the set S_i^k is recursively defined as follows:*
$$S_i^k := \left\{ s_i \in S_i^{k-1} \mid \nexists\, \delta_i \in \mathcal{H}(S_i^{k-1}) \right.$$
such that $u_i(\delta_i, s_{-i}) > u_i(s_i, s_{-i}), \forall\, s_{-i} \in S_{-i}^{k-1} \big\}$.
And define $S_i^\infty := \prod_{k=1}^\infty S_i^k$. The set S_i^∞ is the set of player i's actions that survives iterative elimination of strictly dominated actions.

Observe that after a finite numbers of steps the process of iterated elimination of strictly dominated actions will certainly halt in the sense that there is no action that can be further eliminated, since we restrict our attention to finite games. Moreover, one can show that the procedure of iterated elimination of strictly dominated actions does not depend on the order that we proceed the elimination, that is, it always yields the same surviving set of actions for each player.

With the aid of this procedure, we can thus easily identify the set of rationalizable actions for each player in finite games, which thus provides a nice algorithm for finding rationalizable actions.

Proposition 2.3 (Pearce, 1984). *For any finite normal game G, the set of profiles of rationalizable actions coincides with the set of profiles that survives the process of iterated elimination of strictly dominated actions.*

3 Rationalizability with Imprecise Probabilities

Following the tradition of decision making under uncertainty, the concept of rationalizability assumes that each player's belief regarding the other players' strategies is represented by *a single personal probability measure*. However, there are many convincing arguments for supporting imprecision in beliefs - even in the context of single-agent decision problems (see [Ellsberg, 1961] and [Walley, 1999]). A number of alternative models to subjective expected utility theory have been proposed, which advocate the use of

imprecise probabilities for dealing with uncertainty in decision problems (see, for instance, [Gilboa and Schmeidler, 1989] and [Levi, 1974]). It is thus natural to incorporate these ideas into the traditional game-theoretic framework. Based on the rules of Γ-maximin [Berger, 1985; Gilboa and Schmeidler 1989] and *E-admissibility*, we present here a generalized game-theoretic framework as an initial attempt to examine how modeling uncertainty with imprecise probabilities may provide insight into traditional game theory. In analogy with the concept of rationalizability, we propose two new game-theoretic solution concepts: the solution concept that we shall call Γ-*maximin rationalizability* attempts to capture the idea that each player models the other players as Γ-maximin decision makers, and the other one named *E-rationalizability* is meant to represent the idea that each player thinks of the other players as rational decision makers who respects the *E*-admissibility criterion.

An immediate question that is crucial to this investigation is: which model of imprecise probabilities should be assumed as representation of players' beliefs in strategic situations? There are a variety of mathematical models proposed in the literature to represent uncertainty in single-agent decision problems. For instance, lower previsions, upper and lower probabilities, sets of probabilities, non-additive probabilities, and belief functions (see [Walley, 1991]). Among these widely-discussed models of imprecise probabilities, a plausible method is to use *a convex set of probability distributions*, also called a credal set [Levi, 1980], to represent a decision maker's beliefs when confronted with uncertainty. A great advantage of this approach is that it allows us to deal with any state of insufficiencies in our information, including complete ignorance, in a unified way. Here we adopt this representation of uncertainty as the intended model for the players' beliefs regarding the other players' strategy choices. In order to distinguish it from the previous way of modeling beliefs, we will hereafter refer to a belief as a conjecture. Slightly modifying the formulation of belief in the framework of rationalizability, we define a conjecture of a player as follows:

Definition 7. *In a strategic form game G, a conjecture of player i, denoted by C_{-i}, about the other players' strategy choices is a (nonempty) convex set of probability measures over the opponents' actions S_{-i}.*

Note that this way of representing players' beliefs is a natural generalization of using a single probability distribution, as discussed earlier in the context of rationalizability. Moreover, this representation of beliefs admits the possibility of a correlated conjecture in the sense that, a player's conjecture may contain a probability distribution that cannot be obtained by

independent mixtures over her opponents' strategies, for the elements of a conjecture are probability measures defined over S_{-i}.

One can interpret each member in a player's conjecture as the frequency of the strategy choices by her opponents, each of which is randomly drawn from a large population. More precisely, each player thinks that each of her opponents stands for a large set of players and has the same set of feasible choice. In this context, the probability distributions in player i's conjecture are viewed as the frequencies with which the members of the set S_{-i} are used by those large populations. In light of this, a probability distribution in a conjecture of a player has a completely different meaning from a mixed strategy, even though they may look the same from a mathematical point of view.

Under the preceding interpretation, it is reasonable to consider the cases where the set of strategies for some player is not convex, but players' conjectures are required to be convex. We understand that it is standard practice in game theory to consider the mixed extensions of games, that is, to include all the mixed strategies. Nevertheless, we may want to model circumstances where only the pure strategies are available to the players, which can be suitably described in the current framework with our interpretation.

In the context of single-agent decision making, several decision rules such as Γ-*maximin* [Berger, 1985; Gilboa and Schmeidler, 1989], *E-admissibility* [Levi, 1974], and *maximality* [Walley, 1999] have been discussed in the literature of imprecise probabilities (for a detailed comparison between these criteria see [Schervish et al., 2003], [Seidenfeld, 2004], and [Troffaes, 2007]). There is, however, no general agreement among decision theorists as to which is the right rule for judging rational decisions when uncertainty is expressed by a convex set of probability functions. Among these suggested criteria, the rule of Γ-maximin generalizes the principle of maximizing expected utility by simply taking the lower expected utility, thereby inducing a complete order on the decision set. More precisely, according to Γ-maximin, a rational decision maker should choose an option to maximize the minimum expected value with respect to a convex set of probabilities. This rule for decision making under uncertainty seems suitable for describing decision makers who are *uncertainty averse*, as it always takes the worst possible expected value as the base for maximization. Nevertheless, it has already been noted in [Seidenfeld, 2004] that the rule of Γ-maximin fails to distinguish between open and closed, convex and non-convex sets of probabilities, since choices based on this decision rule essentially reduces to binary comparisons which share the same supporting hyperplanes. It thus implies that the properties of closure and convexity concerning players' conjectures regarding their opponents' strategy choices are indistinguishable by Γ-maximin rationalizability.

The other decision criterion that we shall discuss below is often called *E-admissibility*, which was implicitly mentioned in [Savage, 1954] and extensively advocated by Issac Levi [1974]. According to this decision rule, an option is *E*-admissible if it maximizes expected utility relative to some probability distribution in the convex set of probabilities. In contrast with Γ-maximin, *E*-admissibility does not generate an order of options, but it does avoid the above-mentioned limitation, since it cannot be characterized by pairwise comparisons. As shown in the context of decision making, these two rules are not equivalent in the sense that they may recommend different sets of admissible options. Thus it is not surprising that the game-theoretic solution concepts defined based on these rules are not equivalent either, as illustrated by an example in the next section.

Under strategic situations, players are usually assumed to be uncertain about the other players' strategy behavior, and can only attempt to deduce their opponents' rational actions from the structure of the game and available information about their opponents' preferences. In most games, it is impossible for players to ascertain their opponents' actual behavior. Due to the insufficient information about preferences and irreducible strategic considerations, any level of uncertainty revealed by the imprecision in the set of probabilities may occur in situations of strategic interaction. Since Γ-maximin and *E*-admissibility have been often discussed in the literature of decision theory, it is therefore interesting to study the cases where all the players would use the rule Γ-maximin or *E*-admissibility to choose their strategies in games. By analogy to the framework of rationalizability, we need to be explicit about what we mean by a strategy being rational under uncertainty.

Definition 8. *In a strategic form game G, a strategy δ_i of player i is Γ-rational under uncertainty if there exists a conjecture C_{-i} of player i such that δ_i maximizes player i's minimum expected utility with respect to C_{-i}. In this case, we say that δ_i is a Γ-maximin admissible strategy relative to the conjecture C_{-i}.*

Likewise, we can define a notion called *E*-admissible strategy in a game where players are assumed to use *E*-admissibility as the criterion for strategy choices.

Definition 9. *In a strategic form game G, a strategy δ_i of player i is E-rational under uncertainty if there exists a conjecture C_{-i} of player i such that δ_i maximizes player i's expected utility for some probability*

in C_{-i}. In this case, we say that δ_i is an E-admissible strategy relative to the conjecture C_{-i}.

Recall that the key idea of the concept of rationalizability is that each player regards the other other players as expected utility maximizers. It requires not only that players are rational in the sense of maximizing expected utility with respect some belief, but also that players' beliefs should be consistent with their opponents being rational in a similar way. The solution concept introduced below extends this idea to contexts, where each player is assumed to model the other players as decision makers who employ Γ-maximin or E-admissibility as the decision rule with respect to uncertainty. More specifically, we present a new solution concept that is meant to capture the idea that players are required to consider only those strategies that are rational under uncertainty, and that are supported by conjectures that do not contradict with their opponents being rational under uncertainty.

Now we need to specify the condition for a player's conjecture being consistent with her opponents' rationality in the senses of Definition 8 and Definition 9 rather than in a traditional decision-theoretic sense. A natural suggestion is to require that each element of the conjecture assigns positive probability only to those actions of her opponents that are rational under uncertainty. Putting these ideas together, we can formally define the new solution concept called Γ-*maximin rationalizability.*

Definition 10. *In a strategic form game G, an action $s_i \in S_i$ of player i is Γ-maximin rationalizable if for each player $j \in I$, there exists a set $A_j \subseteq S_j$ of actions such that: (i) $s_i \in A_i$, and (ii) every action s_j in A_j is Γ-maximin admissible relative to some conjecture C_{-j} of player j such that the support of each element of C_{-j} is a subset of A_{-j}.*

According to the above definition, one only needs to find a set of acts and a conjecture for each player in order to check whether a strategy is Γ-maximin rationalizable or not. Unlike the above formulation, Klibanoff [1996] has provided an alternative characterization of rationalizability with uncertainty aversion (see the definition before Theorem 4), which is defined as an iterative reduction process on the strategies. We shall see that his definition turns out to be equivalent to the concept of Γ-maximin rationalizability defined above. As noted in [Osborne, 2004], there are two distinct ways of defining rationalizability: one depends upon an iterated elimination procedure and the other does not. In the light of this, it seems fair to say that Klibanoff's characterization and the above formulation follow exactly the two different ways to generalize rationalizability in normal form games to

accommodate uncertainty aversion, although they actually correspond to the same solution concept.

Analogously, the other solution concept that we call E-rationalizability can be formally defined as follows.

Definition 11. *In a strategic form game G, an action $s_i \in S_i$ of player i is E-rationalizable if for each player $j \in I$, there exists a set $A_j \subseteq S_j$ of actions such that: (i) $s_i \in A_i$, and (ii) every action s_j in A_j is E-admissible relative to some conjecture C_{-j} of player j such that the support of each element of C_{-j} is a subset of A_{-j}.*

4 Discussion of Properties

The aim of this section is to establish some properties of the solution concepts Γ-maximin rationalizability and E-rationalizability. Among other things, we will see that, Γ-maximin rationalizability can reasonably embrace a broader class of strategy profiles as outcomes under certain circumstances in comparison with rationalizability, whereas E-rationalizability can be distinguished from Γ-maximin rationalizability based on the ideas originated in decision theory. In addition, we will characterize the condition under which these three solution concepts coincide.

4.1 General Results

As we have noted, both of the decision rules, Γ-maximin and E-admissibility, can be regarded as simple extensions of the principle of maximizing expected utility to contexts where uncertainty is modeled by a set of probability measures. It is obvious that the former two rules lead to the same recommendations as the latter one when the set of probability measures is a singleton set. This enables us to show that the concepts of Γ-maximin rationalizability and E-admissibility generalize the notion of rationalizability to contexts where a set of probabilities is employed to represent uncertainty in games.

Proposition 4.1. *For any strategic form game G and each player i, if an action s_i^* of player i is rationalizable, then it is Γ-maximin rationalizable. This holds for E-rationalizability as well.*

Proof. Suppose that $s_i^* \in S_i$ is rationalizable. According to Definition 3, we have that there exists a set Z_j of actions for each player $j \in I$ such that both conditions specified in the definition are satisfied. Set $A_j \equiv Z_j$ for every player j. It immediately follows that $s_i^* \in A_i$. And it is clear that every action in A_j is both Γ-maximin admissible and E-admissible relative to some conjecture of player j by considering the set containing only one probability distribution over A_{-j},

as in this case both Γ-maximin and E-admissibility are equivalent to the principle of expected utility maximization. We can thus conclude that s_i^* is Γ-maximin rationalizable, and E-rationalizable as well. □

According to Proposition 2.1, the set of rationalizable actions of each player is nonempty for any finite normal form games. By applying this result, we can easily establish the existences of Γ-maximin rationalizable and E-rationalizable action in strategic games.

Corollary 4.2. *For any strategic form game, there always exists at least one Γ-maximin rationalizable action for each player i. This holds for E-rationalizable action as well.*

4.2 Comparisons

At this point, the reader may wonder whether the sets of Γ-maximin rationalizable and E-rationalizable actions are in fact identical to the set of rationalizable actions. It has already noted in [Epstein, 1997] that the concepts of Γ-maximin rationalizability and rationalizability are not equivalent when the analysis is restricted to only pure strategies. He also includes a generic game (see the game of Figure 1 in [Epstein, 1997]) that is designed to illustrate that difference. Yet he offers no explicit demonstration.

It has been pointed out in [Seidenfeld, 2004] that an option that is Γ-maximin admissible may not be Bayes admissible. Inspired by this result, I show by the following example that the concept of Γ-maximin may induce a larger set of solutions compared to rationalizability. It also serves the purpose of illustrating the definition of Γ-maximin rationalizability.

Example 1. Consider the 3×2 game shown in Figure 1. Unlike the usual setting which includes mixed strategies, we assume here that both players' feasible options are pure strategies only, that is, explicit randomization is excluded; no non-trivial mixed strategy is available to any player.

	L	R
U	$10, 1$	$0, 2$
M	$4, 10$	$4, 1$
D	$0, 1$	$10, 2$

Figure 1: A normal form game

It is easy to verify that only the pure strategies D and R are rationalizable for player 1 and 2 respectively. The previous argument basically relies on the fact that player 1's action M is strictly dominated when mixed strategies are taken into account. As a matter of fact, in this game the set of rationalizable action is the same, regardless of whether we allow explicit randomization or not. To see this, note that the action M is a never-best response, and thus does not belong to the support of any belief of her opponent. Therefore, the restriction imposed on the feasible options of the players does not alter the set of rationalizable actions for both players.

Nevertheless, I claim that all the actions of both player are Γ-maximin rationalizable in the sense of Definition 10. The crucial part for establishing the claim is to see that the action M of player 1 is actually Γ-maximin rationalizable, even though it is not rationalizable. This can be shown by considering the following construction: (i) let the sets of actions for both players be specified as follows: $A_1 = \{U, M\}$ and $A_2 = \{L, R\}$, and (ii) assume that player 1's and player 2's conjecture is depicted respectively by the following convex sets: $C_{-1} = \big\{\mathbb{P}_1(\cdot) : \{L, R\} \to [0, 1] \mid \mathbb{P}_1(\cdot) \text{ is a probability and } 0.2 \leqslant \mathbb{P}_1(R) \leqslant 0.6\big\}$ and $C_{-2} = \big\{\mathbb{P}_2(\cdot) : \{U, M, D\} \to [0, 1] \mid \mathbb{P}_2(\cdot) \text{ is a probability}, \mathbb{P}_2(D) = 0, \text{ and } 0.45 \leqslant \mathbb{P}_2(U) \leqslant 0.95\big\}$.

Under the specifications above, it is obvious that the first condition in Definition 10 is directly satisfied, since the action M belongs to the set A_1 specified for player 1. And it can be seen from Figure 2 and Figure 3 that the second condition is also satisfied, since player 1's lower expected payoff given by the actions U and M is the same with respect to the set C_{-1}, and the actions L and R also yield the same lower expected payoff to player 2 with respect to the set C_{-2}. We can thus say that every action in A_1 and A_2 is Γ-maximin admissible relative to the conjectures C_{-1} and C_{-2} respectively. In addition, note that every probability distribution in C_{-1} and C_{-2} assigns positive probability only to those action in A_2 and A_1 respectively. We can therefore conclude that the action M is Γ-maximin rationalizable. Once M can be Γ-maximin rationalized, it is then straightforward to verify that the other actions of both players are Γ-maximin rationalizable as well.


Figure 2: Expected utility to player 1

This example illustrates that the set of Γ-maximin rationalizable actions may differ from the set of rationalizable actions in some cases. In particular, the former

Figure 3: Expected utility to player 2

solution concept admits the action M as a candidate for the outcome of the game, which is ruled out by the concept of rationalizability. Intuitively, if player 1 is completely ignorant about player 2's strategy choices, it seems quite reasonable for player 1 to select M, as it has the highest security level. Thus one may say that the concept of Γ-maximin rationalizability does capture our intuition in some games.

The result suggested by the above example is not surprising, since the concept of Γ-maximin rationalizability in fact employs a richer representation of uncertainty than that assumed by rationalizability. More precisely, Γ-maximin rationalizability allows each player to model her opponents as Γ-maximin decision makers under uncertainty, which in fact includes the expected utility model considered by rationalizability as a special case. Hence, the concept of Γ-maximin rationalizability gives rise to a broader class of solutions under certain circumstances.

Nevertheless, it has been shown in [Seidenfeld, 2004] that the E-admissibility criterion differs from the rule of Γ-maximin in the context of individual decision making. It is therefore natural to expect that the solution concepts E-rationalizability and Γ-maximin rationalizability would not be equivalent in the game-theoretic context. In order to see this, consider again the game in Example 1. It is easy to see that player 1's option M is not E-admissible for any probability distribution over L, R. Based on this fact, we can then conclude that M is not E-rationalizable, which is Γ-maximin rationalizable as established above. Therefore, E-rationalizability and Γ-maximin rationalizability are not equivalent to each other in the sense that they may lead to different sets of admissible actions for players. It is worthwhile pointing out that E-rationalizability prescribes the same set of admissible actions as the one recommended by rationalizability in this example. It is not difficult to show that this holds for all finite normal form games. In this sense, the concept of E-rationalizability has a more intimate relationship with the traditional notion of rationalizability compared to Γ-maximin rationalizability.

Furthermore, there is another subtle difference be-

tween E-rationalizability and Γ-maximin rationalizability, which is based on some idea in decision theory. As mentioned before, in the context of individual decision making, Γ-maximin fails to distinguish among different convex sets of probabilities, while E-admissibility is capable of distinguishing between any two closed convex sets of probabilities. Putting this into a game-theoretic context, we can show that E-rationalizability and Γ-maximin rationalizability may lead to different sets of admissible options for a player given the same conjecture about opponents' strategy choices. In other words, even though the player holds the same belief model of the other players, E-rationalizability may recommend a different set of admissible options from the other suggested by Γ-maximin rationalizability. To see this, consider Example 1 again. Suppose that player 1's belief about player 2's strategy choice is represented by the conjecture $C_{-1} = \{\mathbb{P}_1(\cdot) : \{L, R\} \to [0, 1] \mid \mathbb{P}_1(\cdot) \text{ is a probability and } 0.4 < \mathbb{P}_1(R) \leqslant 0.6\}$. Under this belief model, both M and D have the same infimum of expectation, and thus they are Γ-maximin admissible. However, only D is E-admissible, since D strictly dominates M with respect to C_{-1}. In this case, E-rationalizability and Γ-maximin rationalizability give rather different recommendations to player 1.

So far, we have shown how the notion of imprecise probabilities sheds light on the traditional game-theoretic framework, by illustrating the difference between Γ-maximin rationalizability and rationalizability, and further by examining the distinction between E-rationalizability and Γ-maximin rationalizability. However, it is also interesting to investigate when these solution concepts turn out to be equivalent. In other words, we want to give the conditions under which the decision rules Γ-maximin and E-admissibility reduce to the principle of expected utility maximization, including in cases where a convex set of probabilities is used to represent uncertainty.

Some basic notation and definitions are necessary for the following discussion. We are concerned here with finite decision problems where uncertainty is modeled by a closed convex set of probability functions. We let Ω denote a finite state space and let $\mathcal{O}$ denote a finite set of outcomes. An option (or act) f is a mapping from the state space Ω to the set of outcomes $\mathcal{O}$. Let $\mathcal{A}$ be a set of options available to the decision maker. As before, we will use the notation $\mathcal{H}(\mathcal{A})$ to denote the convex hull of $\mathcal{A}$. For sake of simplicity, we assume that the decision maker's values for outcomes are determinate and are represented by a cardinal utility function.

Definition 12. *Let $\mathcal{A}$ be a set of options and let $\mathcal{P}$ be a convex set of probability distributions on the un-*

derlying state space Ω. *An option* $f \in \mathcal{A}$ *is* Bayes admissible *with respect to* $\mathcal{P}$ *if there exists* $\mathbb{P} \in \mathcal{P}$ *such that* f *maximizes the expected utility under* $\mathbb{P}$, *that is,* $\mathbb{E}_{\mathbb{P}}(f) \geqslant \mathbb{E}_{\mathbb{P}}(g)$ *for all* $g \in \mathcal{A}$.

The above criterion recommends selecting those options in $\mathcal{A}$ that maximizes expected utility for at least one $\mathbb{P} \in \mathcal{P}$, which corresponds exactly to the idea of *E-admissibility*. We can now present the classic result (see Corollary 3.9.6 in [Walley, 1999] and Theorem 1 in [Schervish et al., 2003]) in decision theory, which plays a crucial role in establishing the central result of this section.

Proposition 4.3. *If the option set* $\mathcal{A}$ *is convex, then every option that is maximal admissible with respect to a closed convex set* $\mathcal{P}$ *of probability distributions is Bayes admissible with respect to* $\mathcal{P}$. *That is, if* $f \in \mathcal{A}$ *is not Bayes admissible, then there exists some* $g \in \mathcal{A}$ *different from* f *such that* $\mathbb{E}_{\mathbb{P}}(g) > \mathbb{E}_{\mathbb{P}}(f)$ *for all* $\mathbb{P} \in \mathcal{P}$.

We can now characterize the condition under which the concepts of Γ-maximin rationalizability and *E*-rationalizability are equivalent to rationalizability.

Proposition 4.4. *For any strategic form game G, if each player's choice set is convex and each player's conjecture regarding her opponents' choices is represented by a closed convex set of probabilities, then the set of Γ-maximin rationalizable actions is equal to the set of rationalizable actions. This holds for E-rationalizability as well.*

Proof. ($\Leftarrow$): It follows directly from Proposition 4.1.

($\Rightarrow$): Consider an arbitrary player $i \in I$. Suppose that s_i is not rationalizable. Then it follows from Proposition 2.3 that s_i is strictly dominated, which, by Lemma 2.2, implies that s_i is a never-best response. It thus follows that s_i is not a Bayes admissible action, since it is not a best response to any belief of player i. Note that each player's choice set is assumed to be convex. Hence, by Proposition 4.3, we have that s_i is not maximal admissible, that is, there exists some δ_i in player i's choice set such that player i's expected payoff to δ_i is strictly greater than her expected payoff to s_i with respect to any correlated belief regarding the other players' strategic behaviors. Accordingly, s_i is not Γ-maximin admissible relative to any conjecture, as any conjecture of player i is a subset of the set of correlated beliefs about her opponents' strategy choices. We can therefore conclude that the action s_i is not Γ-maximin rationalizable, as required.

The result concerning *E*-rationalizability can be established in a similar fashion. $\square$

Klibanoff [1996] also establishes the equivalence between Γ-maximin (or uncertainty aversion) rationaliz-

ability and iterated strict dominance (see Theorem 4), whose proof depends heavily on the equivalence of the iterative definitions of uncertainty aversion rationalizability and rationalizability. By contrast, the proof I present here uses essentially Proposition 4.3, and thus has a decision-theoretic flavor. To some extent, the above proof makes explicit why such an equivalence holds by providing an alternative justification based on an important result in decision theory.

The above result implies that Γ-maximin rationalizability, *E*-admissibility and rationalizability suggest the same set of strategies for each player as rational decisions for games where players are allowed to consider the convex extensions of their choice sets. And it is quite standard in game theory to examine all the mixtures of the pure strategies. In view of this, we may say that the current framework provides a more general theoretical foundation for the concept of rationalizability. That is, the solutions suggested by rationalizability can be supported by a more general decision theory based on weaker assumptions. In that sense, rationalizability is a quite robust solution concept, which is implied merely by the assumption of common knowledge of players being Γ-maximin rational or *E*-rational.

5 Concluding Remarks

A variety of mathematical models have been discussed in the literature to deal with decision making under uncertainty in single-agent decision problems. In contrast with canonical Bayesian decision theory, which uses just one probability function to represent a decision maker's uncertainty, these models use imprecise probabilities, such as a nontrivial set of probability functions, to represent uncertainty. Based on this idea, I have developed in this paper a general theoretical framework for analyzing how different decision rules can be incorporated into the framework of normal-form rationalizability when uncertainty is represented by imprecise probabilities.

More precisely, I extended the notion of rationalizability to the case where players' conjectures about opponents' strategy choices are represented by a convex set of probability measures, instead of a unique probability function. In the spirit of rationalizability, I introduced a solution concept called Γ-*maximin rationalizability*, which captures the idea that each player models the other players as Γ-maximin decision makers with respect to sets of probabilities representing uncertainty; similarly, I also defined another solution concept named *E-rationalizability*. It is easy to see that both Γ-maximin rationalizability and *E-rationalizability* include the concept of rationalizabil-

ity as a special case when the set of probability measures contains only a single probability function. In addition, I have shown by an example that these concepts are not equivalent. I have also identified the conditions under which these solution concepts coincide with each other.

Now I sketch some suggestions for future work along the current line of research. One natural project is to apply some other decision rules like maximality to interactive situations, in a way similar to the framework developed in this paper. And it also seems natural to extend the current framework to the context of extensive form games in which sequential decisions are involved. In this way, one can develop a general theory of games under uncertainty.

References

[1] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, Springer-Verlag, 1985.

[2] D. Bernheim, Rationalizable Strategic Behavior. *Econometrica* 52: 1007-1028, 1984.

[3] J. Dow and S. Werlang, Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction. *J. Econ. Theory* 64: 305-324, 1994.

[4] J. Eichberger and D. Kelsey, Non-Additive Beliefs and Strategic Equilibria. *Games and Economic Behavior* 30: 183âĂŞ215, 2000.

[5] D. Ellsberg, Risk, Ambiguity and the Savage Axioms. *Quart. J. Econom.* 75: 643-669, 1961.

[6] L. Epstein, Preference, Rationalizability and Equilibrium, *J. Econom. Theory* 73: 1-29, 1997.

[7] I. Gilboa and D. Schmeidler, Maxmin Expected Utility with Non-Unique Prior. *J. Math. Econom.* 18: 141-153, 1989.

[8] J. C. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA, 1988.

[9] P. Klibanoff, Uncertainty, Decision and Normal Form Games, mimeo, 1994.

[10] K. C. Lo, Equilibrium in Beliefs under Uncertainty. J. Econom. Theory 71: 443-484, 1996.

[11] I. Levi, On Indeterminate Probabilities. *J. Phil.* 71: 391-418, 1974.

[12] I. Levi, *The Enterprise of Knowledge*. MIT Press, Cambridge, MA, 1980.

[13] H. Liu, Robust Equilibria under Linear Tracing Procedure. In *Proc. of 7th Int. Symp. on Imprecise Probabilities: Theories and Applications*, F. Coolen, G. de Cooman, T. Fetz, and M. Oberguggenberger (eds.), Innsbruck, Austria, 257-266, 2011.

[14] M. J. Osborne, *An Introduction to Game Theory*. Oxford University Press, New York, Oxford, 2004.

[15] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994.

[16] D. Pearce, Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* 52: 1029-1050, 1984.

[17] L. J. Savage, *The Foundations of Statistics*. Wiley, New York, 1954.

[18] T. Seidenfeld, A contrast between two decision rules for use with (convex) sets of probabilities. *Synthese* 140: 69-88, 2004.

[19] M. J. Schervish, T. Seidenfeld, J. B. Kadane, and I. Levi, Extensions of Expected Utility Theory and Some Limitations of Pairwise Comparisons. In *Proc. of 3rd Int. Symp. on Imprecise Probabilities and Their Applications*, J. M. Bernard, T. Seidenfeld, M. Zaffalon (eds.), Lugano, Switzerland, Carleton Scientific, Waterloo, 496-510, 2003.

[20] D. Schmeidler, Subjective Probability and Expected Utility Without Additivity. *Econometrica* 57: 571-587, 1989.

[21] M. Troffaes, Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning* 45: 17-29, 2007.

[22] P. Walley, *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, New York, 1991.