Equilibrium in Beliefs Under Uncertainty [1]

Kin Chung Lo

Department of Economics, University of Toronto,

Toronto, Ontario, Canada M5S 1A1

First Version: March, 1994 This Revision: March, 1995

Abstract

Existing equilibrium concepts for games make use of the subjective expected utility model axiomatized by Savage (1954) to represent players' preferences. Accordingly, each player's beliefs about the strategies played by opponents are represented by a probability measure. Motivated by experimental findings such as the Ellsberg Paradox demonstrating that the beliefs of a decision maker may not be representable by a probability measure, this paper generalizes equilibrium concepts for normal form games to allow for the beliefs of each player to be representable by a closed and convex set of probability measures. The implications of this generalization for the strategy choices and welfare of players are studied. *Journal of Economic Literature* Classification Numbers: C72, D81.

Key words: uncertainty, Ellsberg Paradox, multiple priors, Nash Equilibrium, equilibrium in beliefs

[1] I especially thank Professor Larry G. Epstein for pointing out this topic, and for providing supervision and encouragement. I am also grateful to Professors Mike Peters and Shinji Yamashige for valuable discussions and to an associate editor and a referee for helpful comments. Any remaining errors are my own responsibility.

1. INTRODUCTION

Due to its simplicity and tractability, the subjective expected utility model axiomatized by Savage (1954) has been the most important theory in analysing human decision making under uncertainty. In particular, it is almost universally used in game theory. Using the subjective expected utility model to represent players' preferences, a large number of equilibrium concepts have been developed. The most important is *Nash Equilibrium*. Nash (1951) defines an equilibrium of a normal form game to be a (possibly mixed) strategy for each player with the characteristic that it is best for each player not to deviate from the specified strategy given that all the other players also do not deviate. In an alternative interpretation of *Nash Equilibrium* that is particularly relevant to this paper, the specified (mixed) strategy for a player is not regarded as a conscious randomization, but rather as the conjecture (a probability measure) commonly held by all opponents about that player's pure strategy choice. That is, a *Nash Equilibrium* is not an "equilibrium in strategies" but an "equilibrium in beliefs".[2]

On the other hand, the descriptive validity of the subjective expected utility model has been questioned, for example, because of Ellsberg's (1961) famous mind experiment, a version of which follows. Suppose there are two urns. Urn 1 contains 50 red balls and 50 black balls. Urn 2 contains 100 balls. Each ball in urn 2 can be either red or black but the relative proportions are not specified. Consider the following four acts:

f_1	Win \$100 if the ball drawn from urn 1 is black
f_2	Win \$100 if the ball drawn from urn 1 is red
f_3	Win \$100 if the ball drawn from urn 2 is black
f_4	Win \$100 if the ball drawn from urn 2 is red

Ellsberg argues that the typical preferences for the above acts are $f_1 \sim f_2 \succ f_3 \sim f_4$, where the strict preference $f_2 \succ f_3$ reflects an aversion to the "ambiguity" or "Knightian uncertainty" associated with urn 2. Subsequent experimental studies generally support that people are averse to ambiguity.[3] Such aversion contradicts the subjective expected utility model, as is readily demonstrated for the Ellsberg experiment. In fact, it contradicts any model of preference in which underlying beliefs are represented by a probability measure. (Machina and Schmeidler (1992) call such preferences "probabilistically sophisticated".)

[2]The two interpretations of *Nash Equilibrium* should be viewed as complementary rather than mutually exclusive. It is entirely reasonable for one player to play a strictly mixed strategy (even if he does not have a strict incentive to do so) that is uncertain from the point of view of his opponents. For instance, suppose a factory employer only has time to monitor one of his two workers. The employer may use a random device to determine which worker he is going to monitor and it is natural then to suppose that workers do not know for sure what the random device is.

[3]A summary can be found in Camerer and Weber (1992).

The Ellsberg Paradox has motivated generalizations of the subjective expected utility model.[4] In the multiple priors model axiomatized by Gilboa and Schmeidler (1989), the single prior of Savage is replaced by a closed and convex set of probability measures. The decision maker is said to be uncertainty averse if the set is not a singleton. He evaluates an act by computing the minimum expected utility over the probability measures in his set of priors.

Although the Ellsberg Paradox only involves a single decision maker facing an exogenously specified environment, it is natural to think that ambiguity aversion is also common in decision making problems where more than one person is involved. Since (with two exceptions discussed at length below) existing equilibrium notions of games are defined under the assumption that players are subjective expected utility maximizers, deviations from the Savage model to accommodate aversion to uncertainty make it necessary to redefine equilibrium concepts.

This paper generalizes *Nash Equilibrium* and one of its variations in normal form games to allow for the beliefs of each player to be representable by a closed and convex set of probability measures as in the Gilboa-Schmeidler model.[5] The paper then employs the generalized equilibrium concepts to study the effects of uncertainty aversion on strategic interaction in normal form games.

Note that in order to carry out a *ceteris paribus* study of the effects of uncertainty aversion on how a game is played, the solution concept we use for uncertainty averse players should be different from that for Bayesian players only in terms of attitude towards uncertainty. In particular, the solution concepts should share, as far as possible, comparable epistemic conditions. That is, the requirements on what the players should know about each other's beliefs and rationality underlying the new equilibrium concepts should be "similar" to those underlying familiar equilibrium concepts. This point is emphasized throughout the paper and is used to differentiate the equilibrium concepts proposed here from those proposed by Dow and Werlang (1994) and Klibanoff (1993), also in an attempt to accommodate uncertainty aversion in normal form games.

The paper is organized as follows. Section 2 contains a brief review of the multiple priors model and a discussion of how it is adapted to the context of normal form games. Section 3 defines and discusses the equilibrium concepts (generalizing *Nash Equilibrium* and one of its variations) used in this paper. Section 4 makes use of the equilibrium concepts defined in section 3 to investigate how uncertainty aversion affects the strategy choices and welfare of the players. Section 5 identifies

[4] A summary can be found in Camerer and Weber (1992).

[5] Since the Gilboa-Schmeidler model is intuitive and has been applied elsewhere, for example, Dow and Werlang (1992) and Epstein and Wang (1994), it is adopted to carry out the discussion in this paper. However, it will be clear that some of the conclusions apply more generally.

how uncertainty aversion is related to the structure of a game. Section 6 discusses the epistemic conditions of the equilibrium concepts for uncertainty averse players used in this paper and compares them with those underlying the corresponding equilibrium notions for subjective expected utility maximizing players. Section 7 provides a comparison with Dow and Werlang (1994) and Klibanoff (1993). Section 8 argues that the results in previous sections hold even if we drop the particular functional form of the utility function proposed by Gilboa and Schmeidler (1989) but retain some of its basic properties. Some concluding remarks are offered in section 9.

2. PRELIMINARIES

2.1 Multiple Priors Model

In this section, I provide a brief review of the multiple priors model and a discussion of some of its properties that will be relevant in later sections. For any topological space Y, adopt the Borel σ-algebra Σ_Y and denote by $M(Y)$ the set of all probability measures over Y.[6] Adopt the weak* topology on the set of all finitely additive probability measures over (Y, Σ_Y) and the induced topology on subsets. Let X be the space of outcomes. Let (Ω, Σ_Ω) be the space of states of nature and let $\mathcal{F}$ be the set of all bounded measurable functions from Ω to $M(X)$. That is, $\mathcal{F}$ is the set of two-stage, horse-race/roulette-wheel acts, as in Anscombe and Aumann (1963). $f \in \mathcal{F}$ is called a constant act if $f(\omega) = p \quad \forall \omega \in \Omega$; such an act involves (probabilistic) risk but no uncertainty. For notational simplicity, $M(X)$ also denotes the set of constant acts. For $f, g \in \mathcal{F}$ and $\alpha \in [0, 1]$, $\alpha f + (1-\alpha) g \equiv h$ where $h(\omega) = \alpha f(\omega) + (1-\alpha) g(\omega) \quad \forall \omega \in \Omega$.

The primitive $\succeq$ is a preference ordering over acts. Gilboa and Schmeidler (1989) impose a set of axioms on $\succeq$ that are necessary and sufficient for $\succeq$ to be represented by a numerical function having the following structure: there exists an *affine* function $u : M(X) \to \mathbf{R}$ and a unique, nonempty, closed and convex set $\triangle$ of finitely additive probability measures on Ω such that for all $f, g \in \mathcal{F}$,

$$f \succeq g \Leftrightarrow \min_{p \in \triangle} \int u \circ f dp \geq \min_{p \in \triangle} \int u \circ g dp \tag{2.1.1}$$

It is convenient, but in no way essential, to interpret $\triangle$ as "representing the beliefs underlying $\succeq$"; I provide no formal justification for such an interpretation.

The difference between the subjective expected utility model and the multiple priors model can be illustrated by a simple example. Suppose $\Omega = \{\omega_1, \omega_2\}$ and $X = \mathbf{R}$. Consider an act

[6] The only exception is that when Y is the space of outcomes X, $M(X)$ denotes the set of all probability measures over X with finite supports.

$f \equiv (f(\omega_1), f(\omega_2))$. If the decision maker is a Bayesian[7] and his beliefs over Ω are represented by a probability measure p, the utility of f is

$$p(\omega_1)u(f(\omega_1)) + p(\omega_2)u(f(\omega_2)).$$

On the other hand, if the decision maker is uncertainty averse with the set of priors

$$\triangle = \{p \in M(\{\omega_1, \omega_2\}) \mid p_l \leq p(\omega_1) \leq p_h \quad \text{with} \quad 0 \leq p_l < p_h \leq 1\},$$

then the utility of f is

$$\begin{cases} p_l u(f(\omega_1)) + (1-p_l)u(f(\omega_2)) & \text{if } u(f(\omega_1)) \geq u(f(\omega_2)) \\ p_h u(f(\omega_1)) + (1-p_h)u(f(\omega_2)) & \text{if } u(f(\omega_1)) \leq u(f(\omega_2)). \end{cases}$$

Note that given any act f with $u(f(\omega_1)) > u(f(\omega_2))$, $(\omega_1, p_l; \omega_2, 1-p_l)$[8] can be interpreted as *local probabilistic beliefs* at f in the following sense: there exists an open neighborhood of f such that for any two acts g and h in the neighborhood,

$$g \succeq h \Leftrightarrow p_l u(g(\omega_1)) + (1-p_l)u(g(\omega_2)) \geq p_l u(h(\omega_1)) + (1-p_l)u(h(\omega_2)).$$

That is, the individual behaves like an expected utility maximizer in that neighborhood with beliefs represented by $(\omega_1, p_l; \omega_2, 1-p_l)$. Similarly, $(\omega_1, p_h; \omega_2, 1-p_h)$ represents the local probabilistic beliefs at f if $u(f(\omega_1)) < u(f(\omega_2))$. Therefore, the decision maker who "consumes" different acts may have different local probability measures at those acts.

There are three issues regarding the multiple priors model that will be relevant when the model is applied to normal form games. The first concerns the decision maker's preference for randomization. According to the multiple priors model, preferences over constant acts, that can be identified with objective lotteries over X, are represented by $u(\cdot)$ and thus conform with the von Neumann Morgenstern model. The preference ordering over the set of all acts is quasiconcave. That is, for any two acts $f, g \in \mathcal{F}$ with $f \sim g$, we have $\alpha f + (1-\alpha)g \succeq f$ for any $\alpha \in (0,1)$. This implies that the decision maker may have a strict incentive to randomize among acts.[9]

The second concerns the notion of null event. Given any preference ordering $\succeq$ over acts, define an event $T \subset \Omega$ to be $\succeq$-null as in Savage (1954): T is *$\succeq$-null* if for any acts f, f', g,

$$\begin{bmatrix} f(\omega) & \text{if} & \omega \in T \\ g(\omega) & \text{if} & \omega \notin T \end{bmatrix} \sim \begin{bmatrix} f'(\omega) & \text{if} & \omega \in T \\ g(\omega) & \text{if} & \omega \notin T \end{bmatrix}.$$

[7]In this paper, Bayesian means subjective expected utility maximizer.

[8]Throughout this paper, I use $(y_1, p_1; \ldots; y_m, p_m)$ to denote the probability measure which attaches probability p_i on y_i. For simplicity, I use y for the degenerate probability measure $(y, 1)$.

[9]See section 7 below for arguments for and against this interpretation.

In words, an event T is $\succeq$-null if the decision maker does not care about payoffs in states belonging to T. This can be interpreted as the decision maker knows (or believes) that T can never happen. If $\succeq$ is expected utility preferences, then T is $\succeq$-null if and only if the decision maker attaches zero probability to T. If $\succeq$ is represented by the multiple priors model, then T is $\succeq$-null if and only if *every* probability measure in $\triangle$ attaches zero probability to T.

Finally, the notion of stochastic independence will also be relevant when the multiple priors model is applied to games having more than two players. Suppose the set of states Ω is a product space $\Omega^1 \times \ldots \times \Omega^N$. In the case of a subjective expected utility maximizer, where beliefs are represented by a probability measure $p \in M(\Omega)$, beliefs are said to be stochastically independent if p is a product measure: $p = \times_{i=1}^{N} m_i$ where $m_i \in M(\Omega^i) \quad \forall i$. In the case of uncertainty aversion, the decision maker's beliefs over Ω are represented by a closed and convex set of probability measures $\triangle$. Let $\text{marg}_{\Omega^i}\triangle$ be the set of marginal probability measures on Ω^i as one varies over all the probability measures in $\triangle$. That is,

$$\text{marg}_{\Omega^i}\triangle \equiv \{m_i \in M(\Omega^i) \mid \exists p \in \triangle \text{ such that } m_i = \text{marg}_{\Omega^i} p\}.$$

Following Gilboa and Schmeidler (1989, p.150-151), say that the decision maker's beliefs are *stochastically independent* if

$$\triangle = \text{closed convex hull of } \{\times_{i=1}^{N} m_i \mid m_i \in \text{marg}_{\Omega^i}\triangle \quad \forall i\}.$$

That is, $\triangle$ is the smallest closed convex set containing all the product measures in $\times_{i=1}^{N}\text{marg}_{\Omega^i}\triangle$.

2.2 Normal Form Games

This section defines n-person normal form games where players' preferences are represented by the multiple priors model. Throughout, the index i varies over $1, \ldots, n$ and the indices j and k are assumed to satisfy $i \neq j \neq k$. Therefore, unless specified otherwise, any statement concerning players i, j and k is intended for all $i = 1, \ldots, n$ and $i \neq j \neq k$. As usual, -i denotes the set of all players other than i. Player i's finite pure strategy space is S_i with typical element s_i. The set of pure strategy profiles is $S \equiv \times_{i=1}^{n} S_i$. The game specifies an outcome function $g_i : S \rightarrow X$ for player i. Since mixed strategies induce lotteries over X, we specify an *affine* function $\hat{u}_i : M(X) \rightarrow \mathbf{R}$ to represent player i's preference ordering over $M(X)$. A set of strategy profiles, outcome functions and utility functions determines a normal form game $(S_i, g_i, \hat{u}_i)_{i=1}^{n}$. Let $M(S_i)$ be the set of mixed strategies for player i with typical element σ_i. The set of mixed strategy profiles is therefore given by $\times_{i=1}^{n} M(S_i)$. $\sigma_i(s_i)$ denotes the probability of playing s_i according to the mixed strategy σ_i, $\sigma_{-i}(s_{-i})$ denotes $\prod_{j \neq i} \sigma_j(s_j)$ and σ_{-i} is the corresponding probability measure on S_{-i}.

Assume that player i is uncertain about the strategy profiles of all the other players. Since players (if preferences are strictly quasiconcave) may have a strict incentive to play mixed strategies,[10] the relevant state space for player i is $\times_{j\neq i} M(S_j)$, endowed with the product topology. Each mixed strategy of player i can be regarded as an act over this state space. If player i plays σ_i and the other players play σ_{-i}, i receives the lottery that yields outcome $g_i(s_i, s_{-i})$ with probability $\sigma_i(s_i)\sigma_{-i}(s_{-i})$. Note that this lottery has finite support because S and therefore $\{g_i(s)\}_{s\in S}$ are finite sets. It is also easy to see that the act corresponding to any mixed strategy is bounded and measurable in the sense of the preceding subsection. Consistent with the multiple priors model, player i's beliefs over $\times_{j\neq i} M(S_j)$ are represented by a closed and convex set of probability measures $\hat{B}_i$. Therefore, the objective of player i is to choose $\sigma_i \in M(S_i)$ to maximize

$$\min_{\hat{p}_i\in\hat{B}_i} \int_{\times_{j\neq i} M(S_j)} \sum_{s_i\in S_i} \sum_{s_{-i}\in S_{-i}} \hat{u}_i(g_i(s_i, s_{-i}))\sigma_i(s_i)\sigma_{-i}(s_{-i})d\hat{p}_i(\sigma_{-i}).$$

Define the payoff function $u_i : S \to \mathbf{R}$ as follows: $u_i(s) \equiv \hat{u}_i(g_i(s)) \quad \forall s \in S$. A normal form game can then be denoted alternatively as $(S_i, u_i)_{i=1}^n$ and the objective function of player i can be restated in the form

$$\min_{\hat{p}_i\in\hat{B}_i} \int_{\times_{j\neq i} M(S_j)} \sum_{s_i\in S_i} \sum_{s_{-i}\in S_{-i}} u_i(s_i, s_{-i})\sigma_i(s_i)\sigma_{-i}(s_{-i})d\hat{p}_i(\sigma_{-i}). \tag{2.2.1}$$

In order to produce a simpler formulation of player i's objective function, note that each element in $\hat{B}_i$ is a probability measure over a set of probability measures. Therefore, the standard rule for reducing two-stage lotteries leads to the following construction of $B_i \subseteq M(S_{-i})$:

$$\begin{aligned} B_i \equiv \{p_i \in M(S_{-i}) \mid \exists \hat{p}_i \in \hat{B}_i \\ \text{such that } p_i(s_{-i}) = \int_{\times_{j\neq i} M(S_j)} \sigma_{-i}(s_{-i})d\hat{p}_i(\sigma_{-i}) \quad \forall s_{-i} \in S_{-i}\}. \end{aligned}$$

The objective function of player i can now be rewritten as

$$\min_{p_i\in B_i} u_i(\sigma_i, p_i), \tag{2.2.2}$$

where $\min_{p_i\in B_i} u_i(\sigma_i, p_i) \equiv \min_{p_i\in B_i} \sum_{s_i\in S_i} \sum_{s_{-i}\in S_{-i}} u_i(s_i, s_{-i})\sigma_i(s_i)p_i(s_{-i})$.

Convexity of $\hat{B}_i$ implies that B_i is also convex. Further, from the perspective of the multiple priors model (2.1.1), (2.2.2) admits a natural interpretation whereby S_{-i} is the set of states of

[10] See section 7 below for arguments regarding this approach. Also note that uncertainty aversion is not the only reason for players to have a strict incentive to randomize. In Crawford (1990) and Dekel, Safra and Segal (1991), players may also strictly prefer to randomize even though they are probabilistically sophisticated.

nature relevant to i and B_i is his set of priors over S_{-i}. Because of the greater simplicity of (2.2.2), the equilibrium concepts used in this paper will be expressed in terms of (2.2.2) and B_i instead of (2.2.1) and $\hat{B}_i$. The above construction shows that doing this is without loss of generality. However, the reader should always bear in mind that the former is derived from the latter and I will occasionally go back to the primitive level to interpret the equilibrium concepts.

3. EQUILIBRIUM CONCEPTS

3.1 Equilibrium Concepts for Bayesian Players

This section defines equilibrium concepts for Bayesian players, that is, for subjective expected utility maximizers. The definition of equilibrium proposed by Nash (1951) can be stated as follows:

Definition 1. A *Nash Equilibrium* is a mixed strategy profile $\{\sigma_i^*\}_{i=1}^n$ such that

$$\sigma_i^* \in BR_i(\sigma_{-i}^*) \equiv \operatorname*{argmax}_{\sigma_i \in M(S_i)} u_i(\sigma_i, \sigma_{-i}^*).$$

Under the assumption that players are expected utility maximizers, Nash proves that any finite matrix game of complete information has a *Nash Equilibrium.* It is well known that there are two interpretations of *Nash Equilibrium.* The traditional interpretation is that σ_i^* is the actual strategy used by player i. In a *Nash Equilibrium*, it is best for player i to use σ_i^* given that other players choose σ_{-i}^*. The second interpretation is that σ_i^* is not necessarily the actual strategy used by player i. Instead it represents the marginal beliefs of player j about what pure strategy player i is going to pick. Under this interpretation, *Nash Equilibrium* is usually stated as an n-tuple of probability measures $\{\sigma_i^*\}_{i=1}^n$ such that

$$s_i \in BR_i(\sigma_{-i}^*) \qquad \forall s_i \in \text{support of } \sigma_i^*.$$

Its justification is that given that player i's beliefs are represented by σ_{-i}^*, $BR_i(\sigma_{-i}^*)$ is the set of strategies that maximizes the utility of player i. So player j should "think", if j knows i's beliefs, that only strategies in $BR_i(\sigma_{-i}^*)$ will be chosen by i. Therefore the event that player i will choose a strategy which is not in $BR_i(\sigma_{-i}^*)$ should be "null" (in the sense as defined in section 2.1) from the point of view of player j. This is the reason for imposing the requirement that every strategy s_i in the support of σ_i^*, which represents the marginal beliefs of player j, must be an element of $BR_i(\sigma_{-i}^*)$.

The "beliefs" interpretation of *Nash Equilibrium* allows us to see clearly the source of restrictiveness of this solution concept. First, the marginal beliefs of player j and player k about what player i is going to do are represented by the same probability measure σ_i^*. Second, player i's beliefs about what his opponents are going to do are required to be stochastically independent in the sense that the probability distribution σ_{-i}^* on the strategy choices of the other players is a product measure. We are therefore led to consider the following variation.

Definition 2. A *Bayesian Beliefs Equilibrium* is an n-tuple of probability measures $\{b_i\}_{i=1}^n$ where $b_i \in M(S_{-i})$ such that[11]

$$\text{marg}_{S_i} b_j \in BR_i(b_i) \equiv \underset{\sigma_i \in M(S_i)}{\text{argmax}}\, u_i(\sigma_i, b_i).$$

It is easy to see that if $\{\sigma_i^*\}_{i=1}^n$ is a *Nash Equilibrium*, then $\{\sigma_{-i}^*\}_{i=1}^n$ is a *Bayesian Beliefs Equilibrium*. Conversely, we say that a *Bayesian Beliefs Equilibrium* $\{b_i\}_{i=1}^n$ constitutes a *Nash Equilibrium* $\{\sigma_i^*\}_{i=1}^n$ if $b_i = \sigma_{-i}^*$. Note that in games involving only two players, the two equilibrium concepts are equivalent in that a *Bayesian Beliefs Equilibrium* must constitute a *Nash Equilibrium*.

However, when a game involves more than two players, the definition of *Bayesian Beliefs Equilibrium* is more general. For instance, in a *Bayesian Beliefs Equilibrium*, players i and k can disagree about what player j is going to do. That is, it is allowed that $\text{marg}_{S_j} b_i \neq \text{marg}_{S_j} b_k$.

Example 1: Marginal Beliefs Disagree

Suppose the game involves three players. Player 1 only has one strategy $\{X\}$. Player 2 only has one strategy $\{Y\}$. Player 3 has two pure strategies $\{L, R\}$. The payoff to player 3 is a constant. $\{b_1 = YL, b_2 = XR, b_3 = XY\}$ is a *Bayesian Beliefs Equilibrium*.However it does not constitute a *Nash Equilibrium* because players 1 and 2 disagree about what player 3 is going to do. ∎

Second, in a *Bayesian Beliefs Equilibrium*, player i is allowed to believe that the other players are playing in a correlated manner. As argued by Aumann (1987), this does not mean that the other players are actually coordinating with each other. It may simply reflect that i believes that there exist some common factors among the players that affect their behaviour; for example, player i knows that all other players are professors of economics.

Example 2: Stochastically Dependent Beliefs

Suppose the game involves three players. Player 1 has two pure strategies $\{U, D\}$. Player 2 has two pure strategies $\{L, R\}$. Player 3 has two pure strategies $\{T, B\}$. The payoffs of players 1 and

[11] To avoid confusion, note that this is not Harsanyi's *Bayesian Equilibrium* for games of incomplete information with Bayesian players.

2 are constant. The payoff matrix for player 3 is as follows.[12]

Payoff Matrix for Player 3

	UL	UR	DL	DR
T	-10	3	4	-10
B	0	0	0	0

$$b_1 = (LT, 0.5; RT, 0.5)$$

$$b_2 = (UT, 0.5; DT, 0.5)$$

$$b_3 = (UR, 0.5; DL, 0.5)$$

It is easy to see that $\{b_1, b_2, b_3\}$ listed above is a *Bayesian Beliefs Equilibrium*. Moreover the marginal beliefs of the players agree. However it does not constitute a *Nash Equilibrium*. The reason is that player 3's beliefs about the strategies of players 1 and 2 are stochastically dependent. If player 3 believes that the strategies of player 1 and player 2 are stochastically independent, player 3's beliefs are possibly $(UL, 0.25; UR, 0.25; DL, 0.25; DR, 0.25)$ and T would no longer be his best response. ∎

3.2 Equilibrium Concepts for Uncertainty Averse Players

This section defines generalized equilibrium notions that allow for players' preferences to be represented by the multiple priors model. The generalization of *Bayesian Beliefs Equilibrium* is presented first.

Definition 3. A *Beliefs Equilibrium* is an n-tuple of sets of probability measures $\{B_i\}_{i=1}^n$ where $B_i \subseteq M(S_{-i})$ is a nonempty, closed and convex set such that

$$\text{marg}_{S_i} B_j \subseteq BR_i(B_i) \equiv \underset{\sigma_i \in M(S_i)}{\text{argmax}} \min_{p_i \in B_i} u_i(\sigma_i, p_i).$$

When expressed in terms of $\hat{B}_i$, a *Beliefs Equilibrium* is an n-tuple of closed and convex sets of probability measures $\{\hat{B}_i\}_{i=1}^n$ such that

$$\sigma_i \in BR_i(\hat{B}_i) \qquad \forall \sigma_i \in \cup_{\hat{p}_j \in \hat{B}_j} \text{support of } \text{marg}_{M(S_i)} \hat{p}_j$$

[12] For all n-person games presented in this paper, the payoff is in terms of utility.

where $BR_i(\hat{B}_i)$ is the set of strategies which maximizes (2.2.1).[13]

The interpretation of *Beliefs Equilibrium* parallels that of its Bayesian counterpart. Given that player i's beliefs are represented by $\hat{B}_i$, $BR_i(\hat{B}_i)$ is the set of strategies that maximizes the utility of player i. So player j should "think", if j knows i's beliefs, that only strategies in $BR_i(\hat{B}_i)$ will be chosen by i. Therefore the event that player i will choose a strategy that is not in $BR_i(\hat{B}_i)$ should be "null" (in the sense as defined in section 2.1) from the point of view of player j. This is the reason for imposing the requirement that every strategy σ_i in the union of the support of every probability measure in $\text{marg}_{M(S_i)}\hat{B}_j$, which represents the marginal beliefs of player j about what player i is going to do, must be an element of $BR_i(\hat{B}_i)$.

It is obvious that every *Bayesian Beliefs Equilibrium* is a *Beliefs Equilibrium*. Say that a *Beliefs Equilibrium* $\{B_i\}_{i=1}^n$ is *proper* if not every B_i is a singleton.

Recall that *Nash Equilibrium* is different from *Bayesian Beliefs Equilibrium* in two respects: (i) The marginal beliefs of the players agree and (ii) the overall beliefs of each player are stochastically independent. An appropriate generalization of *Nash Equilibrium* to allow for uncertainty aversion should also possess these two properties. Consider therefore the following definition.

Definition 4. A *Beliefs Equilibrium* $\{B_i\}_{i=1}^n$ is called a *Beliefs Equilibrium with Agreement* if there exists $\times_{i=1}^n \Sigma_i \subseteq \times_{i=1}^n M(S_i)$ such that B_i = closed convex hull of $\{\sigma_{-i} \in M(S_{-i}) \mid \text{marg}_{S_j}\sigma_{-i} \in \Sigma_j\}$.

We can see as follows that this definition delivers the two properties "agreement" and "stochastic independence of beliefs": as explained in section 2.2, player i's beliefs are represented by a convex set of probability measures $\hat{B}_i$ on $\times_{j\neq i} M(S_j)$. I require the marginal beliefs of the players to agree in the sense that $\text{marg}_{M(S_j)}\hat{B}_i = \text{marg}_{M(S_j)}\hat{B}_k$. To capture the idea that the beliefs of each player are stochastically independent, I impose the requirement that $\hat{B}_i$ contains all the product measures. That is,

$$\hat{B}_i = \text{closed convex hull of } \{\times_{j\neq i}\hat{m}_j \mid \hat{m}_j \in \text{marg}_{M(S_j)}\hat{B}_i\}.$$

B_i is derived from $\hat{B}_i$ as in section 2.2. By construction, we have $\text{marg}_{S_j} B_i = \text{marg}_{S_j} B_k$ = convex hull of Σ_j and B_i takes the form required in the definition of *Beliefs Equilibrium with Agreement*.

[13] To be even more precise, a *Beliefs Equilibrium* is an n-tuple of closed and convex sets of probability measures $\{\hat{B}_i\}_{i=1}^n$ such that the complement of $BR_i(\hat{B}_i)$ is a set of $\text{marg}_{M(S_i)}\hat{p}_j$-measure zero for every $\hat{p}_j \in \hat{B}_j$.

Note that *Beliefs Equilibrium* and *Beliefs Equilibrium with Agreement* coincide in two-person games. Further, for n-person games, if $\{b_i\}_{i=1}^n$ is a *Bayesian Beliefs Equilibrium with Agreement*, then $\{b_i\}_{i=1}^n$ constitutes a *Nash Equilibrium.*

To provide further perspective and motivation for the equilibrium concepts just proposed, I state two variations of *Beliefs Equilibrium* and explain why they are *not* the focus of this paper. Given that any strategy in $BR_i(B_i)$ is equally good for player i, it is reasonable for player j to feel completely ignorant about which strategy i will pick from $BR_i(B_i)$. This leads us to consider the following strengthening of *Beliefs Equilibrium*:

Definition 5. A *Strict Beliefs Equilibrium* is a *Beliefs Equilibrium* with $\text{marg}_{S_i} B_j = BR_i(B_i)$.

A *Beliefs Equilibrium* may not be a *Strict Beliefs Equilibrium*, as demonstrated in the following example.[14] The example also shows that a *Strict Beliefs Equilibrium* does not always exist, which is obviously a serious deficiency of this solution concept.

Example 3: Nonexistence of Strict Beliefs Equilibrium

	L	R
U	3,2	-1,2
D	0,4	0,-100

This game only has one *Nash Equilibrium* $\{U, L\}$. It is easy to check that it is not a *Strict Beliefs Equilibrium.* In fact, there is no *Strict Beliefs Equilibrium* for this game. ∎

An opposite direction is to consider weakening the definition of *Beliefs Equilibrium.*

Definition 6. A *Weak Beliefs Equilibrium* is an n-tuple of beliefs $\{B_i\}_{i=1}^n$ such that $\text{marg}_{S_i} B_j \cap BR_i(B_i) \neq \emptyset$.

It is clear that any *Beliefs Equilibrium* is a *Weak Beliefs Equilibrium.* The converse is not true. If $\text{marg}_{S_i} B_j \not\subseteq BR_i(B_i)$, there are some strategies (in j's beliefs about i) that player i will definitely not choose. However, player j considers those strategies "possible". On the other hand, $\text{marg}_{S_i} B_j \cap BR_i(B_i) \neq \emptyset$ captures the idea that player j cannot be "too wrong". It is also not the focus of this paper because we do not expect much strategic interaction if the players know so little about their opponents.[15]

[14] A parallel statement for Bayesian players is that a *Nash Equilibrium* may not be a *Strict Nash Equilibrium.*

[15] *Weak Beliefs Equilibrium* will be discussed further in section 6 (Proposition 8) and section 7, where its relation to the equilibrium concepts proposed by Dow and Werlang (1994) and Klibanoff (1993) is discussed.

Finally it is useful to clarify the relationship between the equilibrium notions defined in this section and some familiar concepts in the received theory of normal form games.

Definition 7. The strategy σ_i^* is a *maximin strategy* for player i if

$$\sigma_i^* \in \underset{\sigma_i \in M(S_i)}{\text{argmax}} \min_{p_i \in M(S_{-i})} u_i(\sigma_i, p_i).$$

The following result is immediate:

Proposition 1 *If $\{M(S_{-i})\}_{i=1}^n$ is a Beliefs Equilibrium, then every $\sigma_i \in M(S_i)$ is a maximin strategy.*

Definition 8. Set $\Gamma_i^0 \equiv M(S_i)$ and recursively define[16]

$$\begin{aligned}\Gamma_i^n = \{\sigma_i \in \Gamma_i^{n-1} \mid \exists p \in M(\times_{j \neq i} \text{supp}\Gamma_j^{n-1})\\ \text{such that } u_i(\sigma_i, p) \geq u_i(\sigma_i', p) \quad \forall \sigma_i' \in \Gamma_i^{n-1}\}.\end{aligned}$$

For player i, the set of *Correlated Rationalizable Strategies* is $\mathcal{R}_i \equiv \cap_{n=0}^{\infty} \Gamma_i^n$. We call $\mathcal{RB}_i \equiv \cap_{n=1}^{\infty} M(\times_{j \neq i} \text{supp}\Gamma_j^{n-1})$ the set of *Rationalizable Beliefs.*

These notions are related to *Beliefs Equilibrium* by the next proposition.

Proposition 2 *Suppose $\{B_i\}_{i=1}^n$ is a Beliefs Equilibrium. Then $BR_i(B_i) \subseteq \mathcal{R}_i$ and $B_i \subseteq \mathcal{RB}_i$.*

Proof:

Set $\hat{\Gamma}_i^0 \equiv M(S_i)$ and recursively define

$$\begin{aligned}\hat{\Gamma}_i^n = \{\sigma_i \in \hat{\Gamma}_i^{n-1} \mid \exists \mathcal{P} \subseteq M(\times_{j \neq i} \text{supp}\hat{\Gamma}_j^{n-1})\\ \text{such that } \min_{p \in \mathcal{P}} u_i(\sigma_i, p) \geq \min_{p \in \mathcal{P}} u_i(\sigma_i', p) \quad \forall \sigma_i' \in \hat{\Gamma}_i^{n-1}\}.\end{aligned}$$

By definition, $\Gamma_i^0 = \hat{\Gamma}_i^0$. It is obvious that $\Gamma_i^1 \subseteq \hat{\Gamma}_i^1$. Any element σ_i not in Γ_i^1 does not survive the first round of the iteration in the definition of correlated rationalizability. Since correlated rationalizability and iterated strict dominance coincide (see Fudenberg and Tirole (1991, p.52)),

[16] The notation suppΓ_j^{n-1} stands for the union of the supports of the probability measures in Γ_j^{n-1}.

there must exist $\sigma_i^* \in \Gamma_i^0$ such that $u_i(\sigma_i^*, p) > u_i(\sigma_i, p) \quad \forall p \in M(\times_{j\neq i}\text{supp}\Gamma_j^0)$. This implies $\min_{p\in\mathcal{P}} u_i(\sigma_i^*, p) > \min_{p\in\mathcal{P}} u_i(\sigma_i, p) \quad \forall\mathcal{P} \subseteq M(\times_{j\neq i}\text{supp}\hat{\Gamma}_j^0)$. Therefore $\sigma_i \notin \hat{\Gamma}_i^1$ and we have $\Gamma_i^1 = \hat{\Gamma}_i^1$. The argument can be repeated to establish $\Gamma_i^n = \hat{\Gamma}_i^n \quad \forall n$.

$BR_i(B_i)$ is rationalized by B_i, that is, $BR_i(B_i) \subseteq \hat{\Gamma}_i^1$. According to the definition of *Beliefs Equilibrium*, $\text{marg}_{S_i} B_j \subseteq BR_i(B_i) \subseteq \hat{\Gamma}_i^1$. This implies $\times_{j\neq i}\text{marg}_{S_j} B_i \subseteq \times_{j\neq i}\hat{\Gamma}_j^1$ and therefore $B_i \subseteq M(\times_{j\neq i}\text{supp}\hat{\Gamma}_j^1)$. The argument can be repeated to establish $BR_i(B_i) \subseteq \hat{\Gamma}_i^n$ and $B_i \subseteq M(\times_{j\neq i}\text{supp}\hat{\Gamma}_j^n) \quad \forall n$. ∎

4. Does Uncertainty Aversion Matter?

4.1 Questions

In section 3, I have set up a framework that enables us to investigate how uncertainty aversion affects strategic interaction in the context of normal form games. My objective here is to address the following two specific questions:

1. As an outside observer, one only observes the actual strategy choice but not the beliefs of each player. Is it possible for an outside observer to distinguish uncertainty averse players from Bayesian players?
2. Does uncertainty aversion make the players worse off (better off)?

To deepen our understanding, let me first provide the answers to the above two questions in the context of single person decision making and conjecture the possibility of extending them to the context of normal form games.

4.2 Single Person Decision Making

The first question is: As an outside observer, can we distinguish an uncertainty averse decision maker from a Bayesian decision maker? The answer is obviously yes if we have "enough" observations. (Otherwise the Ellsberg Paradox would not even exist!) However, it is easy to see that if we only observe an uncertainty averse decision maker who chooses one act from a convex constraint set $\mathcal{G} \subseteq \mathcal{F}$, then his choice can always be rationalized (as long as monotonicity is not violated) by a subjective expected utility function. For example, take the simple case where $\Omega = \{\omega_1, \omega_2\}$. The feasible set of utility payoffs $C \equiv \{(u(f(\omega_1)), u(f(\omega_2))) \mid f \in \mathcal{G}\}$ generated by $\mathcal{G}$ will be a convex set in $\mathbf{R}^2$. Suppose the decision maker chooses a point $c \in C$. To rationalize his choice by an expected utility function, we can simply draw a linear indifference curve which is tangent to C at c, with slope describing the probabilistic beliefs of the expected utility maximizer.

The above answer is at least partly relevant to the first question posed in section 4.1. That is because in a normal form game, an outside observer only observes that each player i chooses a strategy from the set $M(S_i)$. An important difference, though, is that the strategy chosen by i is a best response given his beliefs and these are part of an equilibrium. Therefore it is possible that the consistency condition imposed by the equilibrium concept can enable us to break the observational equivalence. Moreover it should also be expected that the stronger the equilibrium concept, the more likely that uncertainty averse players can be identified.

The second question addresses the welfare consequences of uncertainty aversion: Does uncertainty aversion make a decision maker worse off (better off)? There is a sense in which uncertainty aversion makes a decision maker worse off. For simplicity, suppose again that $X = \mathbf{R}$. Suppose that initially, beliefs over the state space Ω are represented by a probability measure $\hat{p}$ and next that beliefs change from $\hat{p}$ to the set of priors $\triangle$ with $\hat{p} \in \triangle$. Given $f \in \mathcal{F}$, let $CE_{\triangle}(f)$ be the certainty equivalent of f, that is, $u(CE_{\triangle}(f)) = \min_{p \in \triangle} \int u \circ f dp$. Similar meaning is given to $CE_{\hat{p}}(f)$. Then uncertainty aversion makes the decision maker worse off in the sense that

$$CE_{\hat{p}}(f) \geq CE_{\triangle}(f).$$

That is, the certainty equivalent of any f when beliefs are represented by $\hat{p}$ is higher than that when beliefs are represented by $\triangle$.

Note that in the above welfare comparison, I am fixing the utility function of lotteries u. This assumption can be clarified by the following restatement: Assume that the decision maker has a fixed preference ordering $\succeq^*$ over $M(X)$ which satisfies the independence axiom and is represented numerically by u. Denote by $\succeq$ and $\succeq'$ the orderings over acts corresponding to the priors $\hat{p}$ and $\triangle$ respectively. Then the above welfare comparison presumes that both $\succeq$ and $\succeq'$ agree with $\succeq^*$ on the set of constant acts, that is, for any $f, g \in \mathcal{F}$ with $f(\omega) = p$ and $g(\omega) = q$ for all $\omega \in \Omega$, $f \succeq g \Leftrightarrow f \succeq' g \Leftrightarrow p \succeq^* q$.

At this point, it is not clear that the above discussion extends to the context of normal form games. When strategic considerations are present, one might wonder whether it is possible that if player 1 is uncertainty averse and if player 2 knows that player 1 is uncertainty averse, then the behaviour of player 2 is affected in a fashion that benefits player 1 relative to a situation where

2 knows that 1 is a Bayesian.[17] When both players are uncertainty averse and they know that their opponents are uncertainty averse, can they choose a strategy profile that Pareto dominates equilibria generated when players are Bayesians?

4.3 Every Beliefs Equilibrium Contains a Bayesian Beliefs Equilibrium

In this section, the two questions posed in section 4.1 are addressed using the equilibrium concepts *Bayesian Beliefs Equilibrium* and *Beliefs Equilibrium*. The answers are contained in the following proposition.

Proposition 3 *If $\{B_i\}_{i=1}^n$ is a Beliefs Equilibrium, then there exist $b_i \in B_i$, $i = 1, \ldots, n$, such that $\{b_i\}_{i=1}^n$ is a Bayesian Beliefs Equilibrium. Moreover, $BR_i(B_i) \subseteq BR_i(b_i)$.*

Proof: [18]

It is sufficient to show that there exists $b_i \in B_i$ such that $BR_i(B_i) \subseteq BR_i(b_i)$. This and the fact that $\{B_i\}_{i=1}^n$ is a *Beliefs Equilibrium* imply

$$\text{marg}_{S_i} b_j \in \text{marg}_{S_i} B_j \subseteq BR_i(B_i) \subseteq BR_i(b_i).$$

Therefore $\{b_i\}_{i=1}^n$ is a *Bayesian Beliefs Equilibrium.*

We have that $u_i(\cdot, p_i)$ is linear on $M(S_i)$ for each p_i and $u_i(\sigma_i, \cdot)$ is linear on B_i for each σ_i. Therefore, by Fan's Theorem (Fan (1953)),

$$\underline{u} \equiv \max_{\sigma_i \in M(S_i)} \min_{p_i \in B_i} u_i(\sigma_i, p_i) = \min_{p_i \in B_i} \max_{\sigma_i \in M(S_i)} u_i(\sigma_i, p_i).$$

By definition, $\sigma_i \in BR_i(B_i)$ if and only if $\min_{p_i \in B_i} u_i(\sigma_i, p_i) = \underline{u}$. Therefore,

$$u_i(\sigma_i, p_i) \geq \underline{u} \quad \forall p_i \in B_i \quad \forall \sigma_i \in BR_i(B_i). \tag{4.3.1}$$

Take $b_i \in \text{argmin}_{p_i \in B_i} \max_{\sigma_i \in M(S_i)} u_i(\sigma_i, p_i)$. Then conclude that

$$u_i(\sigma_i, b_i) \leq \underline{u} = \max_{\sigma_i \in M(S_i)} u_i(\sigma_i, b_i) \quad \forall \sigma_i \in M(S_i). \tag{4.3.2}$$

[17] A well known example where this kind of reasoning applies is the following. An expected utility maximizer who is facing an exogenously specified set of states of nature, always prefers to have more information before making a decision. However, this is not necessarily the case if the decision maker is playing a game against another player. The reason is that if player 1 chooses to have less information and if player 2 "knows" it, the strategic behaviour of player 2 may be affected. The end result is that player 1 may obtain a higher utility by throwing away information. (See the discussion of correlated equilibrium in chapter 2 of Fudenberg and Tirole (1991).)

[18] Though I prove a result (Proposition 12) below for more general preferences, I provide a separate proof here because the special structure of the Gilboa-Schmeidler model permits some simplification.

Combining (4.3.1) and (4.3.2), we have

$$u_i(\sigma_i, b_i) = \underline{u} \quad \forall \sigma_i \in BR_i(B_i), \text{ that is, } BR_i(B_i) \subseteq BR_i(b_i).$$

∎

Example 4[19]: Illustrating Proposition 3

	C_1	C_2	C_3	C_4
R_1	0,1	0,1	1,1	0,1
R_2	0,1	0,1	0,1	1,1
R_3	1,1	-1,1	0,1	0,1
R_4	-1,1	1,1	0,1	0,1

$$B_1 = M(\{C_1, C_2, C_3, C_4\})$$

$$B_2 = \{R_1\}$$

$$BR_1(B_1) = \{p \in M(\{R_1, R_2, R_3, R_4\}) \mid p(R_3) = p(R_4)\}$$

$$BR_2(B_2) = M(\{C_1, C_2, C_3, C_4\})$$

The *Beliefs Equilibrium* $\{B_1, B_2\}$ listed above contains the *Bayesian Beliefs Equilibrium* $\{b_1 = (C_1, 0.5; C_2, 0.5), b_2 = R_1\}$. Note that $BR_1(b_1) = M(\{R_1, R_2, R_3, R_4\})$. This shows that in general, $BR_i(B_i)$ can be a proper subset of $BR_i(b_i)$. The example also demonstrates that a *Proper Beliefs Equilibrium* in general contains more than one *Bayesian Beliefs Equilibrium*. For instance, $\{b_1' = C_3, b_2' = R_1\}$ is another *Bayesian Beliefs Equilibrium*. However $BR_1(b_1') = \{R_1\}$. Therefore not every *Bayesian Beliefs Equilibrium* $\{b_i'\}_{i=1}^n$ contained in a *Beliefs Equilibrium* $\{B_i\}_{i=1}^n$ has the property $BR_i(B_i) \subseteq BR_i(b_i')$. ∎

For games involving more than two players, a *Beliefs Equilibrium* in general does not contain a *Nash Equilibrium*. This is already implied by the fact that a *Bayesian Beliefs Equilibrium* is itself a *Beliefs Equilibrium* but not a *Nash Equilibrium*. However, since *Bayesian Beliefs Equilibrium* and *Nash Equilibrium* are equivalent in two-person games, Proposition 3 has the following corollary.

Corollary of Proposition 3 *In a two-person game, if* $\{B_1, B_2\}$ *is a Beliefs Equilibrium, then there exists* $\sigma_j^* \in B_i$ *such that* $\{\sigma_1^*, \sigma_2^*\}$ *is a Nash Equilibrium. Moreover,* $BR_i(B_i) \subseteq BR_i(\sigma_j^*)$.

[19] For all two-person games presented in this paper, player 1 is the row player and player 2 is the column player.

Proposition 3 delivers two messages. The first regards the prediction of how the game will be played. Suppose $\{B_i\}_{i=1}^n$ is a *Beliefs Equilibrium*. The associated prediction regarding strategies played is that i chooses some $\sigma_i \in BR_i(B_i)$. According to Proposition 3, it is always possible to find at least one *Bayesian Beliefs Equilibrium* $\{b_i\}_{i=1}^n$ contained in $\{B_i\}_{i=1}^n$ such that the observed behaviour of the uncertainty averse players (the actual strategies they choose) is consistent with utility maximization given beliefs represented by $\{b_i\}_{i=1}^n$. This implies that an outsider who can only observe the actual strategy choices in the single game under study will not be able to distinguish uncertainty averse players from Bayesian players. (I will provide reasons to qualify such observational equivalence in the next section.)

We can use Proposition 3 also to address the welfare consequences of uncertainty aversion, where the nature of our welfare comparisons is spelled out in section 4.2. If $\{B_i\}_{i=1}^n$ is a *Beliefs Equilibrium*, it contains a *Bayesian Beliefs Equilibrium* $\{b_i\}_{i=1}^n$, and therefore,

$$\max_{\sigma_i \in M(S_i)} \min_{p_i \in B_i} u_i(\sigma_i, p_i) \leq \max_{\sigma_i \in M(S_i)} u_i(\sigma_i, b_i).$$

The left hand side of the above inequality is the ex ante utility of player i when his beliefs are represented by B_i and the right hand side is ex ante utility when beliefs are represented by b_i. The inequality implies that ex ante, i would prefer to play the *Bayesian Beliefs Equilibrium* $\{b_i\}_{i=1}^n$ to the *Beliefs Equilibrium* $\{B_i\}_{i=1}^n$. In that ex ante sense, uncertainty aversion makes the players worse off.[20]

4.4 Uncertainty Aversion can be Beneficial When Players Agree

The comparisons above addressed the effects of uncertainty aversion when the equilibrium concepts used, namely *Beliefs Equilibrium* and *Bayesian Beliefs Equilibrium*, do not require agreement between agents. Here I re-examine the effects of uncertainty aversion when agreement is imposed, as incorporated in the *Beliefs Equilibrium with Agreement* and *Nash Equilibrium* solution concepts.

For two-person games, the Corollary of Proposition 3 still applies since agreement is not an issue given only two players. However, for games involving more than two players, the following example demonstrates that it is possible to have a *Beliefs Equilibrium with Agreement* not containing any *Nash Equilibrium.*

[20] For the *Bayesian Beliefs Equilibrium* $\{b_i\}_{i=1}^n$ constructed in the proof of Proposition 3, we actually have

$$\max_{\sigma_i \in M(S_i)} \min_{p_i \in B_i} u_i(\sigma_i, p_i) = \max_{\sigma_i \in M(S_i)} u_i(\sigma_i, b_i).$$

This, of course, does not exclude the possibility that there may exist other *Bayesian Beliefs Equilibria* contained in $\{B_i\}_{i=1}^n$ such that the equality is replaced by a strict inequality.

Example 5: Uncertainty Aversion Leads to Pareto Improvement

The game presented in this example is a modified version of the prisoners' dilemma. The game involves three players, 1, 2 and $\mathcal{N}$. Player $\mathcal{N}$ can be interpreted as "nature". The payoff of player $\mathcal{N}$ is a constant and his set of pure strategies is $\{X, Y\}$. Players 1 and 2 can be interpreted as two prisoners. The set of pure strategies available for players 1 and 2 are $\{C_1, D_1\}$ and $\{C_2, D_2\}$ respectively, where C stands for "co-operation" and D stands for "defection".

Payoff Matrix for Player 1

	XC_2	YC_2	XD_2	YD_2
C_1	c	c	e	e
D_1	a	b	d	d

Payoff Matrix for Player 2

	XC_1	YC_1	XD_1	YD_1
C_2	c	c	e	e
D_2	b	a	d	d

$$a > c > b \quad \text{and} \quad c > d > e \quad \text{and} \quad 2c < a + b \tag{4.4.1}$$

Note that the game is the prisoner's dilemma game if the inequalities $a > c > b$ in (4.4.1) are replaced by $a = b > c$. (When $a = b > c$, the payoffs of players 1 and 2 for all strategy profiles do not depend on nature's move.) This game is different from the standard prisoners' dilemma in one respect. In the standard prisoners' dilemma game, the expression $a = b > c$ says that it is always better for one player to play D given that his opponent plays C. In this game, the expression $a > c > b$ says that if one player plays D and one plays C, the player who plays D may either gain or lose. The interpretation of the inequalities $c > d > e$ in (4.4.1) is the same as that in the standard prisoners' dilemma. That is, it is better for both players to play C rather than D. However, a player should play D given that his opponent plays D. Note that the last inequality $2c < a + b$ in (4.4.1) is implied by $a = b > c$ in the prisoners' dilemma game. The inequality $2c < a + b$ can be rewritten as $(a - c) - (c - b) > 0$. For player 1, for example, $(a - c)$ is the utility gain from playing D_1 instead of C_1 if the true state is XC_2. $(c - b)$ is the corresponding utility loss if the true state is YC_2. Therefore the interpretation of $2c < a + b$ is that if you know your opponent plays C, the possible gain (loss) for you to play D instead of C is high (low).

Assume that players 1 and 2 know each other's action but they are uncertain about nature's move. To be precise, suppose that the beliefs of the players are

$$B_{\mathcal{N}} = \{C_1 C_2\}$$

$$B_1 = \{p \in M(\{XC_2, YC_2\}) \mid p_l \leq p(XC_2) \leq p_h \quad \text{with} \quad 0 \leq p_l < p_h \leq 1\}$$

$$B_2 = \{p \in M(\{XC_1, YC_1\}) \mid p_l \leq p(XC_1) \leq p_h \quad \text{with} \quad 0 \leq p_l < p_h \leq 1\}$$

The construction of $\{B_{\mathcal{N}}, B_1, B_2\}$ reflects the fact that the players agree. For example, the marginal beliefs of players 1 and 2 regarding $\{X, Y\}$ agree with

$$\triangle = \{p \in M(\{X, Y\}) \mid p_l \leq p(X) \leq p_h \quad \text{with} \quad 0 \leq p_l < p_h \leq 1\}. \tag{4.4.2}$$

Given $\{B_{\mathcal{N}}, B_1, B_2\}$, the payoffs of each pure strategy profile for players 1 and 2 are

	C_2	D_2
C_1	c, c	e, $p_h b + (1 - p_h)a$
D_1	$p_l a + (1 - p_l)b$,e	d, d

Therefore $\{B_{\mathcal{N}}, B_1, B_2\}$ is a *Beliefs Equilibrium with Agreement* and $BR_1(B_1) = \{C_1\}$ and $BR_2(B_2) = \{C_2\}$ if and only if

$$p_l < \frac{c-b}{a-b} \quad \text{and} \quad p_h > \frac{a-c}{a-b}.$$

Note that our assumptions guarantee that $\frac{c-b}{a-b} > 0$ and $\frac{a-c}{a-b} < 1$, so there exist values for p_l and p_h consistent with the above inequalities.

However, $\{B_{\mathcal{N}}, B_1, B_2\}$ does not contain a *Nash Equilibrium*. To see this, suppose that players 1 and 2 are Bayesians who agree that $p(X) = \lambda = 1 - p(Y)$ for some $\lambda \in [0, 1]$. Then they are playing the following game:

	C_2	D_2
C_1	c, c	e, $\lambda b + (1 - \lambda)a$
D_1	$\lambda a + (1 - \lambda)b$,e	d, d

The strategy profile $\{C_1, C_2\}$ is a *Nash Equilibrium* if and only if

$$c \geq \lambda a + (1 - \lambda)b \quad \text{and} \quad c \geq \lambda b + (1 - \lambda)a.$$

There exists $\lambda \in [0, 1]$ such that $\{C_1, C_2\}$ is a *Nash Equilibrium* if and only if

$$c \geq \frac{1}{2}(a + b),$$

which contradicts the last inequality in (4.4.1). Therefore, it is never optimal for both Bayesian players to play C and any *Nash Equilibrium* requires both players to play D and therefore that both players receive d with certainty. In the *Beliefs Equilibrium with Agreement* constructed above, on the other hand, both players play C and receive $c > d$ with certainty.

To better understand why uncertainty aversion leads to a better equilibrium in this game, let us go back to the beliefs $\{B_1, B_2\}$ of players 1 and 2. As explained in section 2.1, although the global beliefs of players 1 and 2 on $\{X, Y\}$ are represented by the same $\triangle$ in (4.4.2), the local probability measures for different acts may be different. For example, the local probability measure on $\{X, Y\}$ at the act corresponding to D_1 is $(X, p_l; Y, 1-p_l)$ and for D_2 it is $(X, p_h; Y, 1-p_h)$ respectively. In the sense of local probability measures, therefore, players 1 and 2 disagree on the relative likelihood of X and Y when they are consuming the acts D_1 and D_2 respectively. This allows playing D to be undesirable for both players.

The example delivers two messages. First, it shows that in a game involving more than two players, uncertainty aversion can lead to an equilibrium that Pareto dominates all *Nash Equilibria*. Second, interpreting player $\mathcal{N}$ in the above game as "nature", the game becomes a two-person game where the players are uncertain about their own payoff functions. Therefore uncertainty aversion can be "beneficial" even in two-person games. ∎

5. Why Do We Have Uncertainty Aversion?

The next question I want to address is: When should we expect (or not) the existence of an equilibrium reflecting uncertainty aversion? In the context of single person decision theory, we do not have much to say about the origin or precise nature of beliefs on the set of states of nature. However, we should be able to say more in the context of game theory. The beliefs of the players should be "endogenous" in the sense of depending on the structure of the game. For example, it is reasonable to predict that the players will not be uncertainty averse if there is an "obvious way" to play the game.

The following two examples identify possible reasons for players to be uncertainty averse.

Example 6: Nonunique Equilibria

	L	C	R
U	1,1	1,1	1 billion,1
M	1,1	1,1	1 billion,1
D	1,1 billion	1,1 billion	1 billion,1 billion

In the above game, any strategy profile is a *Nash Equilibrium*. Any $\{B_1, B_2\}$ is a *Beliefs Equilibrium*. Uncertainty aversion in this game is due to the fact that the players do not have any idea about how their opponents will play. ∎

Example 7: Nonunique Best Responses

	L	R
U	0,1	1,0.5
D	0,1	0,2

In this game, $\{U, L\}$ is the only *Nash Equilibrium*. However, it is equally good for player 1 to play D if he believes that player 2 plays L. Under this circumstance, it may be too demanding to require player 2 to attach probability one to player 1 playing U. At the other extreme, where 2 is totally ignorant of 1's strategy choice, we obtain the *Proper Beliefs Equilibrium* $\{B_1 = \{L\}, B_2 = M(\{U, D\})\}$. ∎

This example shows that the existence of a unique *Nash Equilibrium* is not sufficient to rule out an equilibrium with uncertainty aversion. However, I can prove the following:

Proposition 4 *If the game has a unique Bayesian Beliefs Equilibrium and it is also a strict Nash Equilibrium, then there does not exist a Proper Beliefs Equilibrium.*

Proof:

Let $\{b_i\}_{i=1}^n$ be the unique *Bayesian Beliefs Equilibrium*. Since it is also a strict *Nash Equilibrium* of the game, there exists $s_i^* \in S_i$ such that $b_i = s_{-i}^*$. Let $\{B_i\}_{i=1}^n$ be a *Beliefs Equilibrium*. According to Proposition 3, $s_{-i}^* \in B_i$. Using Proposition 3 and the definition of *Beliefs Equilibrium*, we have $\text{marg}_{S_i} B_j \subseteq BR_i(B_i) \subseteq BR_i(s_{-i}^*) = \{s_i^*\}$. This implies $\{B_i\}_{i=1}^n = \{s_{-i}^*\}_{i=1}^n$. ∎

Corollary of Proposition 4 *In a two-person game, if the game has a unique Nash Equilibrium and it is also a Strict Nash Equilibrium, then there does not exist a Proper Beliefs Equilibrium.*

A *Proper Beliefs Equilibrium* can be ruled out also if the game is dominance solvable.

Proposition 5 *If the game is dominance solvable, then there does not exist a Proper Beliefs Equilibrium.*[21]

[21]We may also want to ask the reverse question: Are the conditions stated in Propositions 4 and 5 necessary for

Proof:

Let $\{B_i\}_{i=1}^n$ be a *Beliefs Equilibrium.* According to Proposition 2, $BR_i(B_i) \subseteq \mathcal{R}_i$ and $B_i \subseteq \mathcal{RB}_i$. Since iterated strict dominance and correlated rationalizability are equivalent, a dominance solvable game has a unique pure strategy profile $\{s_i^*\}_{i=1}^n$ such that $\mathcal{R}_i = s_i^*$ and $\mathcal{RB}_i = s_{-i}^*$. Therefore, $BR_i(B_i) = s_i^*$ and $B_i = s_{-i}^*$. ∎

6. Decision Theoretic Foundation

Recently, decision theoretic foundations for Bayesian solution concepts have been developed. See, for example, Aumann (1987), Aumann and Brandenberger (1991), Brandenberger (1992) and Tan and Werlang (1988). The purpose of this line of research is to understand the knowledge requirements needed to justify equilibrium concepts. Although research on the generalization of equilibrium concepts to allow for uncertainty averse preferences has already started (see section 7 below), serious study of the epistemic conditions for those generalized equilibrium concepts has not yet been carried out.

In this section, I provide epistemic conditions for the equilibrium concepts proposed in this paper. The main finding is that *Beliefs Equilibrium* (*Beliefs Equilibrium with Agreement*) and *Bayesian Beliefs Equilibrium* (*Nash Equilibrium*) presume similar knowledge requirements. **This supports our equilibrium concepts for games with uncertainty aversion and the interpretation of results in previous sections as reflecting solely the effects of uncertainty aversion.** To focus on the intuition, the propositions are only informally discussed in this section. Their proofs can be found in the appendix.

The framework I use to discuss epistemic matters is adopted (with suitable modification) from Aumann (1987).[22] Let Ω be a common finite set of states of nature for the players with typical element ω. H_i is the partition of Ω for player i and $H_i(\omega)$ is the element of H_i that contains ω. The payoff function of player i is $u_i : \Omega \times S \rightarrow \mathbf{R}$. Denote by $f_i : \Omega \rightarrow M(S_i)$ the strategy used by player i so that $f_i(\omega)(s_i)$ is the probability that i plays s_i in state ω. u_i and f_i are adapted

the absence of *Proper Beliefs Equilibrium*? The following game has two *Nash Equilibria* (and therefore it is not dominance solvable). They are $\{U, R\}$ and $\{D, L\}$. However, there does not exist a *Proper Beliefs Equilibrium.*

	L	R
U	2,1	1,2
D	2,2	0,2

[22] In the case where players are expected utility maximizers, all notation in this paragraph will still apply except that $\triangle_i(\omega)$ is a singleton for all ω and f_i is a mapping from Ω to S_i.

to H_i. For each $\omega \in \Omega$, player i's beliefs over $H_i(\omega)$ are represented by a closed and convex set of probability measures $\triangle_i(\omega)$. His beliefs over S_{-i} are represented by a closed and convex set of probability measures $B_i(\omega)$ that is induced from $\triangle_i(\omega)$ in the following way:

$$\begin{aligned} B_i(\omega) \equiv \{ & p_i \in M(S_{-i}) \mid \exists q_i \in \triangle_i(\omega) \\ & \text{such that} \quad p_i(s_{-i}) = \sum_{\hat{\omega} \in H_i(\omega)} q_i(\hat{\omega}) \Pi_{j \neq i} f_j(\hat{\omega})(s_j) \quad \forall s_{-i} \in S_{-i} \}. \end{aligned}$$

The above specification is common knowledge among the players. Player i is said to *know* an event E at ω if $H_i(\omega) \subseteq E$. Say that an event is *mutual knowledge* if everyone knows it. Let $\mathcal{H}$ be the meet of the partitions of all the players and $\mathcal{H}(\omega)$ the element of $\mathcal{H}$ which contains the element ω. An event E is *common knowledge* at ω if and only if $\mathcal{H}(\omega) \subseteq E$. Say that player i is *rational* at ω if his strategy $f_i(\omega)$ maximizes utility as stated in (2.2.2) when beliefs are represented by $B_i(\omega)$.

The following proposition describes formally the knowledge requirements for $\{B_i\}_{i=1}^n$ to be a *Beliefs Equilibrium.* If each B_i is a singleton, then a parallel result for *Bayesian Beliefs Equilibrium* is obtained. (The version of Proposition 6 for *Bayesian Beliefs Equilibrium* for two person-games can be found in Theorem A in Aumann and Brandenberger (1991, p.7).[23])

Proposition 6 *Suppose that at some state ω, the rationality of the players, $\{u_i\}_{i=1}^n$ and $\{B_i\}_{i=1}^n$ are mutual knowledge. Then $\{B_i\}_{i=1}^n$ is a Beliefs Equilibrium.*

The idea of Proposition 6 is not difficult. At ω, player i *knows* j's beliefs $B_j(\omega)$, payoff function $BR_j(\omega)$ and that j is rational. Therefore any strategy $f_j(\omega')$ with $\omega' \in H_i(\omega)$ of player j ,where player i thinks is possible, must be player j's best response given his beliefs. That is, $f_j(\omega') \in BR_j(\omega)(B_j(\omega)) \quad \forall \omega' \in H_i(\omega)$. Since the preference ordering of player j is quasiconcave, any convex combination of strategies in the set $\{f_j(\omega') \mid \omega' \in H_i(\omega)\}$ must also be a best response for player j. By construction, $\text{marg}_{S_j} B_i(\omega)$ is a subset of the convex hull of $\{f_j(\omega') \mid \omega' \in H_i(\omega)\}$. This implies $\text{marg}_{S_j} B_i(\omega) \subseteq BR_j(\omega)(B_j(\omega))$ and therefore $\{B_i\}_{i=1}^n$ is a *Beliefs Equilibrium.*

In a *Beliefs Equilibrium with Agreement*, the beliefs $\{B_i\}_{i=1}^n$ of the players over the strategy choices of opponents are required to have the properties of agreement and stochastic independence. Since B_i is derived from $\triangle_i$, it is to be expected that some restrictions on $\triangle_i$ are needed for $\{B_i\}_{i=1}^n$ to possess the desired properties. In the case where players are expected utility maximizers so that $\triangle_i(\omega)$ is a singleton for all ω, Theorem B in Aumann and Brandenberger (1991, p.8) shows that

[23]In their paper, the notion of knowledge means "ascribe probability one", which is more general than the notion "absolute certainty without possibility of error" that is being used here.

by restricting $\{\triangle_i\}_{i=1}^n$ to come from a common prior, *mutual knowledge* of rationality and payoff functions and *common knowledge* of beliefs are sufficient to imply *Nash Equilibrium.* In the case where players are uncertainty averse, the following proposition says that by restricting each player i to be *completely ignorant* at each ω about the relative likelihood of states in $H_i(\omega)$, exactly the same knowledge requirements that imply *Nash Equilibrium* also imply *Beliefs Equilibrium with Agreement.*

Proposition 7 *Suppose that $\triangle_i(\omega) = M(H_i(\omega)) \quad \forall\omega$. Suppose that at some state ω, the rationality of the players and $\{u_i\}_{i=1}^n$ are mutual knowledge and that $\{B_i\}_{i=1}^n$ is common knowledge. Then $\{B_i\}_{i=1}^n$ is a Beliefs Equilibrium with Agreement.*

The specification of $\triangle_i(\omega)$ as the set of *all* probability measures on $H_i(\omega)$ reflects the fact that player i is *completely ignorant* about the relative likelihood of states in $H_i(\omega)$. Note that Theorem B in Aumann and Brandenberger (1991) is not a special case of Proposition 7. The restriction on $\triangle_i$ imposed by their theorem coincides with that in Proposition 7 only in the case where $H_i(\omega) = \{\omega\} \quad \forall\omega$. Therefore the examples they provide to show the sharpness of their theorem also do not apply to Proposition 7. To serve this purpose, I provide Example 8 to show that Proposition 7 is tight in the sense that *mutual knowledge* rather than *common knowledge* of $\{B_i\}_{i=1}^n$ is not sufficient to guarantee a *Beliefs Equilibrium with Agreement.*

Example 8: Mutual Knowledge of Beliefs is not Sufficient for Agreement

The game consists of three players. The set of states of nature, the information structures and the strategies of the players are listed below.

$$\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$$

$$H_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$$

$$H_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$$

$$H_3 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$$

	ω_1	ω_2	ω_3	ω_4
f_1	L	L	R	R
f_2	U	D	U	D
f_3	T	T	T	T

Suppose $\triangle_i(\omega) = M(H_i(\omega)) \quad \forall\omega$. At ω_1, the beliefs $\{B_i(\omega_1)\}_{i=1}^3$ of the players are *mutual knowledge*. For example, since $B_1(\omega) = M(\{UT, DT\}) \quad \forall\omega \in \Omega$, $B_1(\omega_1)$ is *common knowledge* and therefore *mutual knowledge* at ω_1. According to the proof of Proposition 7, marginal beliefs of the players agree. For example, $\text{marg}_{S_2} B_1(\omega_1) = \text{marg}_{S_2} B_3(\omega_1) = M(\{U, D\})$. However $B_3(\omega_1) = M(\{LU, LD, RU\})$ is not *common knowledge* at ω_1. Player 3 does not *know* that player 1 *knows* player 3's beliefs. At ω_1, player 3 cannot exclude the possibility that the true state is ω_3. At ω_3, player 1 only knows that player 3's beliefs are either represented by $B_3(\omega_3) = M(\{LU, LD, RU\})$ or $B_3(\omega_4) = \{RD\}$. Note that $B_3(\omega_1)$ does not take the form required in the definition of *Beliefs Equilibrium with Agreement.* ∎

Finally, although the notion of *Weak Beliefs Equilibrium* defined in section 3 is not the main focus of this paper, it is closely related to the equilibrium concepts proposed by the papers discussed in section 7 below. According to the following proposition, *complete ignorance* and *rationality* at a state ω are sufficient to imply *Weak Beliefs Equilibrium.*

Proposition 8 *Suppose that at some state ω, $\triangle_i(\omega) = M(H_i(\omega))$ and that players are rational. Then $\{B_i\}_{i=1}^n$ is a Weak Beliefs Equilibrium.*

By looking at the definition of *Weak Beliefs Equilibrium* more carefully, Proposition 8 is hardly surprising. For instance, given any n-person normal form game, it is immediate that $\{M(S_{-i})\}_{i=1}^n$ is a *Weak Beliefs Equilibrium.* Note that Proposition 8 requires only that the players be *rational*; they do not need to *know* that their opponents are *rational.* They also do not need to *know* anything about the beliefs of their opponents.

7. Related Literature

In this section, I compare my equilibrium concepts with those proposed by two other papers in this area.

7.1 Equilibrium Notions

The first is Dow and Werlang (1994). They consider two-person games and assume that players' preference orderings over acts are represented by the convex capacity model proposed by Schmeidler (1989). Any such preference ordering is a member of the multiple priors model (Gilboa and Schmeidler (1989)). The equilibrium concept proposed by Dow and Werlang can be restated using the multiple priors model as follows:

Definition 9. $\{B_1, B_2\}$ is a *Nash Equilibrium Under Uncertainty* if the following conditions are satisfied:

1. There exists $E_i \subseteq S_i$ such that $p_j(T_i) < 1 \quad \forall p_j \in B_j \quad \forall T_i \subset E_i$ and $p_j(E_i) = 1$ for at least one $p_j \in B_j$.
2. $\min_{p_i \in B_i} u_i(s_i, p_i) \geq \min_{p_i \in B_i} u_i(s_i', p_i) \qquad \forall s_i \in E_i \qquad \forall s_i' \in S_i$.

Dow and Werlang (1994) interpret E_i in condition 1 as an event that is infinitely more likely to happen than its complement according to the beliefs B_j of player j. Condition 2 says that every $s_i \in E_i$ is a best response for i, given that B_i represents i's beliefs about the strategy choice of player j.

Another related paper is Klibanoff (1993). He adopts the multiple priors model to represent players' preferences in static normal form games with any finite number of players and proposes the following solution concept:[24]

Definition 10. $(\{\sigma_i^*\}_{i=1}^n, \{B_i\}_{i=1}^n)$ is an *Equilibrium with Uncertainty Aversion* if the following conditions are satisfied:

1. $\sigma_{-i}^* \in B_i$.
2. $\min_{p_i \in B_i} u_i(\sigma_i^*, p_i) \geq \min_{p_i \in B_i} u_i(\sigma_i, p_i) \qquad \forall \sigma_i \in M(S_i)$.

σ_i^* is the actual strategy used by player i and B_i is his beliefs about opponents' strategy choices. Condition 1 says that player i's beliefs cannot be "too wrong". That is, the strategy profile σ_{-i}^* chosen by other players should be considered "possible" by player i. Condition 2 says that σ_i^* is a best response for i given his beliefs B_i.

7.2 Choice of Strategy Space

One important difference (or similarity) among the equilibrium concepts is the choice of strategy space. In Dow and Werlang (1994), the strategy space of each player i is his set of pure strategies S_i. In Klibanoff (1993) and in this paper, the strategy space of each player i is his set of mixed strategies $M(S_i)$. To illustrate the difference, let us consider the following example.[25]

[24] The equilibrium concept presented here is a simplified version. Klibanoff (1993) assumes that players' beliefs are represented by lexicographic sets of probability measures.

[25] This example is taken from Dow and Werlang (1991).

Example 9: Strict Incentive to Randomize

	L	R
U	-2,-2	0,1
D	1,0	-1,-1

$$B_1 = \{p_1 \in M(\{L, R\}) \mid 0.6 \leq p_1(R) \leq 0.9\}$$

$$B_2 = \{p_2 \in M(\{U, D\}) \mid 0.6 \leq p_2(D) \leq 0.9\}$$

In this game, $\{B_1, B_2\}$ is a *Nash Equilibrium Under Uncertainty.* Each player is indifferent between his two pure strategies. However, $\{B_1, B_2\}$ is not a *Beliefs Equilibrium (with Agreement)* as defined in this paper. The reason is that if the utility function of player i over the set of mixed strategies is represented by (2.2.2), then it is strictly better for i to play a strictly mixed strategy given his beliefs B_i. The utility of player 1 derived from the mixed strategy $(U, 0.5; D, 0.5)$ is -0.5, while the utility of each pure strategy is only -0.8. ∎

Example 9 shows that whether a beliefs profile $\{B_1, B_2\}$ constitutes an equilibrium depends on the choice of strategy space. It is therefore important to understand the justification for using one strategy space instead of the other. The justification offered here can be summarized as follows: The use of pure vs. mixed strategy spaces depends on the perception of the players about the order of strategy choices. The adoption of a mixed strategy space in Klibanoff (1993) and in this paper can be justified by the assumption that each player perceives himself as moving last. On the other hand, we can understand the adoption of a pure strategy space in Dow and Werlang (1994) as assuming that each player perceives himself as moving first. (See also Dekel, Safra and Segal (1991) for another instance where the perception of the players about the order of moves is important.)

To see this, let us go back to the context of single person decision theory. Recall that $\succeq$ is a preference ordering over the set of acts $\mathcal{F}$, where each act f maps Ω into $M(X)$. The interpretation of f is as follows. First a horse race determines the true state of nature $\omega \in \Omega$. The decision maker is then given the objective lottery ticket $f(\omega)$. He spins the roulette wheel as specified by $f(\omega)$ to determine the actual prize he is going to receive. Also recall that for any two acts $f, f' \in \mathcal{F}$ and $\alpha \in [0, 1]$, $\alpha f + (1 - \alpha)f'$ refers to the act which yields the lottery ticket $\alpha f(\omega) + (1 - \alpha)f'(\omega)$ in state ω.

Suppose $\succeq$ is quasiconcave (and not linear) as in the case of the Gilboa-Schmeidler model. We can argue as follows that this leads to an incentive to randomize: Suppose a decision maker has to choose between f and f'. Suppose further that he perceives that nature moves first; that is, a particular state $\omega^* \in \Omega$ has been realized but the decision maker does not know what ω^* is. If the decision maker randomizes between choosing f and f' with probabilities α and $1-\alpha$ respectively, he will receive the lottery $\alpha f(\omega) + (1-\alpha)f'(\omega)$ when $\omega^* = \omega$. This is precisely the payoff of the act $\alpha f + (1-\alpha)f'$ in state ω. That is, randomization enables him to enlarge the choice set from $\{f, f'\}$ to $\{\alpha f + (1-\alpha)f' \mid \alpha \in [0,1]\}$. Therefore there will "typically" be an $\alpha \in (0,1)$ such that $\alpha f + (1-\alpha)f'$ is optimal according to $\succeq$.[26]

On the other hand, suppose the decision maker moves first and nature moves second. If the decision maker randomizes between choosing f and f' with probabilities α and $1-\alpha$ respectively, he faces the lottery $(f, \alpha; f', 1-\alpha)$, that delivers act f with probability α and f' with probability $1-\alpha$. Therefore, randomization delivers the set $\{(f, \alpha; f', 1-\alpha) \mid \alpha \in [0,1]\}$ of lotteries over G. Note that $\{(f, \alpha; f', 1-\alpha) \mid \alpha \in [0,1]\}$ is not in the domain of $\mathcal{F}$ and so the Gilboa-Schmeidler model is silent on the decision maker's preference ordering over this set.

The above discussion translates to the context of normal form games with uncertainty averse players as follows. The key is whether player i perceives himself as moving first or last. As explained in section 2.2, each pure strategy of player i can be regarded as an act over the state space $\times_{j \neq i} M(S_j)$. Klibanoff (1993) and I can be viewed as assuming that player i perceives himself as the last person to make the strategy choice. If he randomizes between s_i and s_i' with probabilities α and $1-\alpha$ respectively, he faces the act $\alpha s_i + (1-\alpha)s_i'$. On the other hand, Dow and Werlang (1994) can be viewed as assuming that player i perceives himself as the first person to make the strategy choice. The perception of each player about the order of strategy choices is not observable and there does not seem to be a compelling theoretical case for assuming either order. Therefore, it would seem that either specification of strategy space merits study.

An objection to the present approach that might be raised is that it contradicts "Ellsberg type" behaviour. The argument goes as follows: Suppose a decision maker can choose between f_3 and f_4 listed on p.1. If the decision maker randomizes between f_3 and f_4 with equal probability, it will generate the act $\frac{1}{2}f_3 + \frac{1}{2}f_4$ which yields the lottery $[\$100, \frac{1}{2}; \$0, \frac{1}{2}]$ in each state. Therefore $\frac{1}{2}f_3 + \frac{1}{2}f_4$ is as desirable as f_1 or f_2. This implies that the decision maker will be indifferent between having the opportunity to choose an act from $\{f_1, f_2\}$ or from $\{f_3, f_4\}$ and the Ellsberg

[26] Note that this only explains why the decision maker may strictly prefer to randomize. We also need to rely on the dynamic consistency argument proposed by Machina (1989) to ensure that the decision maker is willing to obey the randomization result after the randomizing device is used. See also Dekel, Safra and Segal (1991, p.241) for discussion of this issue in the context of normal form games.

Paradox disappears! The discussion in previous paragraphs gives us the correct framework to handle this objection. Randomization between f_3 and f_4 with equal probability will generate the act $\frac{1}{2}f_3 + \frac{1}{2}f_4$ only if either the decision maker is explicitly told or he himself perceives that he can first withdraw a ball from the urn but not look at its colour, then toss a fair coin and choose f_3 (f_4) when head (tail) comes up. (Raiffa (1961, p.693)). Also, the preference pattern $f_1 \sim f_2 \succ f_3 \sim f_4$ is already sufficient to constitute one version of the Ellsberg Paradox. In this version, consideration of randomization is irrelevant. Therefore assuming a strict incentive to randomize does not make every version of the Ellsberg Paradox disappear.

7.3 Epistemic Conditions

Another important point of comparison is epistemic foundations. I suggested in the introduction that in order to carry out a *ceteris paribus* study of the effects of uncertainty aversion on how a game is played, we should ensure that the generalized equilibrium concept is different from *Nash Equilibrium* only in one dimension, players' attitude towards uncertainty. In particular, the generalized solution concept should share comparable knowledge requirements with *Nash Equilibrium*. According to this criterion, I argue in section 6 that the solution concepts I propose are appropriate generalizations of their Bayesian counterparts.

Dow and Werlang (1994) and Klibanoff (1993) do not provide epistemic foundations for their solution concepts and a detailed study is beyond the scope of this paper, However, I show below that in the context of two person normal form games, exactly the same epistemic conditions that support *Weak Beliefs Equilibrium* as stated in Proposition 8, namely *complete ignorance* and *rationality*, also support their solution concepts. Therefore the sufficient conditions for players' beliefs to constitute a *Nash Equilibrium Under Uncertainty* or an *Equilibrium with Uncertainty Aversion* do not require them to know anything about the beliefs and rationality of their opponents.

The weak epistemic foundation for their equilibrium concepts is readily reflected by the fact that given any two-person normal form game, $\{M(S_2), M(S_1)\}$ is always a *Nash Equilibrium Under Uncertainty* and there always exist σ_1^* and σ_2^* such that $(\sigma_1^*, \sigma_2^*, M(S_2), M(S_1))$ is an *Equilibrium with Uncertainty Aversion*. The equilibrium notions in these two papers therefore do not fully exploit the difference between a game, where its payoff structure (e.g. dominance solvability) may limit the set of "reasonable" beliefs, and a single person decision making problem, where any set of priors (or single prior in the Bayesian case) is "rational". In fact, Dow and Werlang (1994, p.313) explicitly adopt the view that the degree of uncertainty aversion is subjective, as in the single agent setting, rather than reasonably tied to the structure of the game. As a result, their equilibrium concept delivers a continuum of equilibria for every normal form game (see their theorem on p.313).

Let me now proceed to the formal statements. Since Dow and Werlang (1994) adopt a pure strategy space, I keep all notation from section 6 but redefine

- $f_i : \Omega \to S_i$
- $B_i(\omega) \equiv \{p_i \in M(S_j) \mid \exists q_i \in \triangle_i(\omega)$
  such that $p_i(s_j) = \sum_{\hat{\omega} \in H_i(\omega) \cap \{\omega' | f_j(\omega')=s_j\}} q_i(\hat{\omega}) \quad \forall s_j \in S_j\}$
- $BR_i(B_i) \equiv \mathrm{argmax}_{s_i \in S_i} \min_{p_i \in B_i} u_i(s_i, p_i)$.

That is, $f_i(\omega)$ is the pure action used by player i at ω and $B_i(\omega)$ is the set of probability measures on S_j induced from $\triangle_i(\omega)$ and f_j. It represents the beliefs of player i at ω about j's strategy choice.

Proposition 9 *Suppose that at some state ω, $\triangle_i(\omega) = M(H_i(\omega))$ and that the players are rational. Then $\{B_1, B_2\}$ is a Nash Equilibrium Under Uncertainty.*

Proof: See the appendix. ∎

The demonstration that the same epistemic conditions imply Klibanoff's equilibrium concept is even more straightforward. Note that although Klibanoff's equilibrium notion involves both the specification of beliefs and the actual strategies used by the players while the equilibrium concepts in my paper involve only the former, his equilibrium notion can be readily rewritten as a *Weak Beliefs Equilibrium.* It follows that Proposition 8 provides the epistemic conditions underlying Klibanoff's equilibrium concept as simplified here.

Proposition 10 *$\{B_1, B_2\}$ is a Weak Beliefs Equilibrium if and only if there exists $\sigma_i \in B_j$ such that $(\{\sigma_1, \sigma_2\}, \{B_1, B_2\})$ is an Equilibrium with Uncertainty Aversion.*

Note that Propositions 8 and 9 do not exploit the fact that the information structure of the players are represented by partitions. The two propositions and their proof continue to hold as long as the beliefs of player i at ω are represented by the set of all probability measures over an event $H_i(\omega) \subseteq \Omega$ with the property that $\omega \in H_i(\omega)$. Therefore the conclusion that *complete ignorance* and *rationality* imply the equilibrium concepts of Dow and Werlang (1994) and Klibanoff (1993) remains valid even if we adopt weaker notions of knowledge that are not necessarily defined in terms of partitional information structures.

Another way to see how the epistemic conditions for the equilibrium concepts of Dow and Werlang (1994) and Klibanoff (1993) differ from those underlying my equilibrium concepts is as follows. Assume that each player does not have a strict incentive to randomize. Take S_{-i} to be the state space of player i and S_i to be the set of acts which map S_{-i} to $\mathbf{R}$. Suppose that $\succeq_i$ represents the preference ordering of player i over S_i. Player i is rational if he chooses s_i such that $s_i \succeq_i \hat{s}_i \quad \forall \hat{s}_i \in S_i$. The following is an appropriate restatement of *Nash Equilibrium* in terms of preferences:

Definition 11. $\{\succeq_i, \succeq_j\}$ is a *Nash Equilibrium* if the following conditions are satisfied:

1. There exists $\Gamma_i \subseteq S_i$ such that the complement of Γ_i is $\succeq_j$-null.
2. $s_i \succeq_i \hat{s}_i \qquad \forall s_i \in \Gamma_i \qquad \forall \hat{s}_i \in S_i$.

In words, $\{\succeq_i, \succeq_j\}$ is a *Nash Equilibrium* if the event that player i is irrational is $\succeq_j$-null. *Beliefs Equilibrium (with Agreement)*, when restated in terms of preferences, satisfies Definition 11, with S_i replaced by $M(S_i)$ and using the definition of nullity as stated in section 2.1. The equilibrium concepts of Dow and Werlang (1994) and Klibanoff (1993) satisfy Definition 11 if the notion of nullity in Dow and Werlang (1994) is adopted: an event is $\succeq_j$-null if it is attached zero probability by at least one probability measure in B_j.[27] The adoption of weaker notion of nullity is the reason that in Example 1 of Dow and Werlang (1994, p.314), player 1 chooses a non-rationalizable strategy and in Example 3 of Dow and Werlang (1994, p.316), players do not backward induct in the finite prisoner's dilemma game. The fact that the players are uncertainty averse is not in and of itself the factor that drives those results. It is impossible to deliver such predictions using my equilibrium concept, even if we assume that players do not randomize.

7.4 Agreement and Stochastic Independence

Finally, since both Klibanoff (1993) and I look at games having any finite number of players, a direct comparison of the two equilibrium concepts for n-person normal form games can also be conducted.

It is pointed out in section 7.3 that in terms of the best response property, *Equilibrium with Uncertainty Aversion* only requires $\text{marg}_{S_i} B_j \cap BR_i(B_i) \neq \emptyset$ but *Beliefs Equilibrium with Agreement*

[27]To see that Klibanoff's equilibrium concept satisfies Definition 11 when Dow and Werlang's definition of nullity is adopted, restate *Weak Beliefs Equilibrium* in terms of $\hat{B}_i$: $\{\hat{B}_i, \hat{B}_j\}$ is a *Weak Beliefs Equilibrium* if there exists $\hat{b}_j \in \hat{B}_j$ such that $\sigma_i \in BR_i(\hat{B}_i) \quad \forall \sigma_i \in$ support of $\hat{b}_j$. Set Γ_i in Definition 11 to be the support of $\hat{b}_j$.

requires $\text{marg}_{S_i} B_j \subseteq BR_i(B_i)$. These requirements continue to hold in n-person games.

When we look at games with more than two players, the issues of agreement and stochastic independence become relevant. It is easy to see that condition 1 in Definition 10 only implies

- $\text{marg}_{S_j} B_i \cap \text{marg}_{S_j} B_k \neq \emptyset$
- B_i contains at least one product measure.

Therefore the agreement and stochastic independence properties inherited in an *Equilibrium with Uncertainty Aversion* are much weaker than

- $\text{marg}_{S_j} B_i = \text{marg}_{S_j} B_k$
- B_i contains all the product measures in $\times_{j \neq i} \text{marg}_{S_j} B_i$

which are required to hold in a *Beliefs Equilibrium with Agreement.*

All of above lead to the following proposition:

Proposition 11 *If $\{B_i\}_{i=1}^n$ is a Beliefs Equilibrium with Agreement, then for any $\sigma_i \in \text{marg}_{S_i} B_j$, it is the case that $(\{\sigma_i\}_{i=1}^n, \{B_i\}_{i=1}^n)$ is an Equilibrium with Uncertainty Aversion.*

8. More General Preferences

The purpose of this section is to show that even if we drop the particular functional form proposed by Gilboa and Schmeidler (1989) but retain some of its basic properties, counterparts of the preceding equilibrium concepts and results can be formulated and proven.

Let us first go back to the context of single person decision theory and define a class of utility functions that generalizes the multiple priors model. Recall the notation introduced in section 2.1, whereby $\succeq$ is a preference ordering over the set of acts $\mathcal{F}$, where each act maps Ω into $M(X)$. Impose the following restrictions on $\succeq$: Suppose that $\succeq$ restricted to constant acts conforms to expected utility theory and so is represented by an affine $u : M(X) \rightarrow \mathbf{R}$. Suppose that there exists a nonempty, closed and convex set of probability measures $\triangle$ on Ω such that $\succeq$ is representable by a utility function of the form

$$f \longmapsto U(f) \equiv V(\{\int u \circ f dp \mid p \in \triangle\}) \tag{8.1}$$

for some function $V : \mathbf{R}^{\triangle} \rightarrow \mathbf{R}$. Assume that $\succeq$ is *monotonic* in the sense that for any $f, g \in \mathcal{F}$, if $\int u \circ f dp > \int u \circ g dp \quad \forall p \in \triangle$, then $f \succ g$. Say that $\succeq$ is *multiple priors representable* if $\succeq$ satisfies all the above properties. Quasiconcavity of $\succeq$ will also be imposed occasionally.

Two examples are provided here to clarify the structure of the utility function U in (8.1). Suppose there exists a probability measure μ over $M(\Omega)$ and a concave and increasing function h such that

$$U(f) = \int_{M(\Omega)} h(\int_{\Omega} u \circ f dp) d\mu.$$

In this example, the set of probability measures $\triangle$ corresponds to the support of μ. The interpretation of the above utility function is that the decision maker views an act f as a two-stage lottery. However, the reduction of compound lotteries axiom may not hold (Segal (1990)). Note that this utility function satisfies quasiconcavity. Another example for U which is not necessarily quasiconcave is the Hurwicz (1951) criterion:

$$U(f) = \alpha \min_{p \in \triangle} \int_{\Omega} u \circ f dp + (1 - \alpha) \max_{p \in \triangle} \int_{\Omega} u \circ f dp$$

where $0 \leq \alpha \leq 1$.

Adapting the model to the context of normal form games as in section 2.2, the objective function of player i is $V(\{u_i(\sigma_i, b_i) \mid b_i \in B_i\})$. All equilibrium notions can be defined precisely as before. I now prove the following extension of Proposition 3.

Proposition 12 *Consider a n-person game. Suppose that the preference ordering of each player is multiple priors representable and quasiconcave. If $\{B_i\}_{i=1}^n$ is a Beliefs Equilibrium, then there exists $b_i \in B_i$ such that $\{b_i\}_{i=1}^n$ is a Bayesian Beliefs Equilibrium. Moreover, $BR_i(B_i) \subseteq BR_i(b_i)$.*

Proof:

As in the proof of Proposition 3, it suffices to show that, given B_i, there exists $b_i \in B_i$ such that $BR_i(B_i) \subseteq BR_i(b_i)$.

I first show that for each $\sigma_i \in BR_i(B_i)$, there exists $b_i \in B_i$ such that $\sigma_i \in BR_i(b_i)$. Suppose that this were not true. Then there exists $\sigma_i \in BR_i(B_i)$ such that for each $b_i \in B_i$, we can find $\sigma_i' \in M(S_i)$ with $u_i(\sigma_i, b_i) < u_i(\sigma_i', b_i)$. This implies that there exists $\sigma_i^* \in M(S_i)$ such that $u_i(\sigma_i^*, b_i) > u_i(\sigma_i, b_i) \quad \forall b_i \in B_i$. (See Lemma 3 in Pearce (1984, p.1048).) Since preference of player i is *monotonic*, player i should strictly prefer σ_i^* to σ_i when his beliefs are represented by B_i. This contradicts the fact that $\sigma_i \in BR_i(B_i)$.

Quasiconcavity of preference implies that $BR_i(B_i)$ is a convex set. Therefore there exists an element $\sigma_i \in BR_i(B_i)$ such that the support of σ_i is equal to the union of the support of every probability measure in $BR_i(B_i)$. Since $\sigma_i \in BR_i(b_i)$ for some $b_i \in B_i$ and $u_i(\cdot, b_i)$ is linear on $M(S_i)$, this implies that $s_i \in BR_i(b_i) \quad \forall s_i \in$ support of σ_i. This in turn implies $BR_i(B_i) \subseteq BR_i(b_i)$. ∎

Besides Proposition 3, it is not difficult to see that Proposition 2 also holds if the preference ordering of each player is *multiple priors representable*. (The *monotonicity* of the preference ordering for player i ensures that $\Gamma_i^n = \hat{\Gamma}_i^n \ \forall n$ in the proof of Proposition 2.) Propositions 4 and 5 are also valid because their proof depends only on Propositions 2 and 3. Finally, all the results in section 6, except Proposition 6 which also requires preferences to be quasiconcave, are true under the assumption of *multiple priors representable* preferences.

9. Concluding Remarks

Let me first summarize the questions addressed in this paper:

1. What is a generalization of *Nash Equilibrium* (and its variants) in normal form games to allow for uncertainty averse preferences?
2. What are the epistemic conditions for those generalized equilibrium concepts?
3. As an outside observer, can we distinguish uncertainty averse players from Bayesian players?
4. Does uncertainty aversion make the players worse off (better off)?
5. How is uncertainty aversion related to the structure of the game?

Generalizations of *Nash Equilibrium* have already been proposed by Dow and Werlang (1994) and Klibanoff (1993) to partly answer questions 3, 4 and 5. One important feature of the equilibrium concepts presented in this paper that is different from Dow and Werlang (1994) but common with Klibanoff (1993) is the adoption of mixed instead of pure strategy space. They can both be justified by different perceptions of the players about the order of strategy choices. On the other hand, I can highlight the following relative merits of the approach pursued here. A distinctive feature of the solution concepts proposed in my paper is their epistemic foundation, that resemble as closely as possible those underlying the corresponding Bayesian equilibrium concepts. As pointed out by Dow and Werlang (1994, p.313), their equilibrium concepts are only "presented intuitively rather than derived axiomatically". In my paper, some epistemic conditions are also provided for the equilibrium concepts proposed by Dow and Werlang (1994) and Klibanoff (1993). The weakness of their equilibrium concepts is revealed by the fact that the epistemic conditions do not involve any strategic considerations. This point was demonstrated in section 7.3 where I noted that in any normal form game, regardless of its payoff structure, the beliefs profile $\{M(S_2), M(S_1)\}$ constitutes an equilibrium in their sense.

Using my equilibrium concepts, negative answers are given to questions 3 and 4 for two-person games. However, a three-person game is constructed to demonstrate that there may exist a *Beliefs Equilibrium with Agreement* which Pareto dominates (both ex ante and ex post) every *Nash Equilibrium*. One direction of future research is to identify the general class of economic problems where uncertainty aversion is beneficial to the agents.

Another area of future research concerns generalization of other well known equilibrium concepts for normal form games with Bayesian players. They include, for example, *Correlated Rationalizability* and *Correlated Equilibrium*. If the players have a strict incentive to randomize, the proof of Proposition 2 shows that $\Gamma_i^n = \hat{\Gamma}_i^n \quad \forall n$. That is, generalizing *Correlated Rationalizability* to allow for uncertainty aversion does not make a difference in terms of the set of rationalizable strategies. However, this no longer holds if we adopt the assumption that players do not have a strict incentive to randomize.[28] For *Correlated Equilibrium*, there does not seem to exist a sensible generalization that allows for uncertainty averse preferences.

Appendix

Proof of Proposition 6

Fix ω at the state where the rationality of the players, $\{u_i\}_{i=1}^n$ and $\{B_i\}_{i=1}^n$ are *mutual knowledge*. Player i *knows* player j's beliefs means

$$H_i(\omega) \subseteq \{\omega' \in \Omega \mid B_j(\omega') = B_j(\omega)\}.$$

[28]Consider the following game.

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

Suppose player 1 does not randomize. The strategy D cannot be rationalized by any Bayesian beliefs b_1 because U is better than D if $b_1(L) \geq 0.5$ and M is better than D if $b_1(L) \leq 0.5$. However D can be rationalized by the beliefs $B_1 = M(\{L, R\})$.

Player i *knows* player j is *rational* means

$$H_i(\omega) \subseteq \{\omega' \in \Omega \mid f_j(\omega') \in BR_j(\omega')(B_j(\omega'))\}.$$

Note that BR_j varies with the state because u_j does. Player i *knows* player j's payoff function means

$$H_i(\omega) \subseteq \{\omega' \in \Omega \mid BR_j(\omega') = BR_j(\omega)\}.$$

Therefore

$$\begin{aligned} H_i(\omega) \quad \subseteq \quad & \{\omega' \in \Omega \mid B_j(\omega') = B_j(\omega)\} \\ & \cap \{\omega' \in \Omega \mid f_j(\omega') \in BR_j(\omega')(B_j(\omega'))\} \\ & \cap \{\omega' \in \Omega \mid BR_j(\omega') = BR_j(\omega)\} \\ \subseteq \quad & \{\omega' \in \Omega \mid f_j(\omega') \in BR_j(\omega)(B_j(\omega))\}. \end{aligned}$$

This implies

$$\{f_j(\omega') \mid \omega' \in H_i(\omega)\} \subseteq BR_j(\omega)(B_j(\omega)).$$

The fact that the preference of player j is quasiconcave implies that $BR_j(\omega)(B_j(\omega))$ is a convex set. Therefore we have

$$\text{convex hull of } \{f_j(\omega') \mid \omega' \in H_i(\omega)\} \subseteq BR_j(\omega)(B_j(\omega)).$$

By construction of $B_i(\omega)$,

$$\text{marg}_{S_j} B_i(\omega) \subseteq \text{convex hull of } \{f_j(\omega') \mid \omega' \in H_i(\omega)\} \subseteq BR_j(\omega)(B_j(\omega)).$$

This shows that $\{B_i\}_{i=1}^n$ is a *Beliefs Equilibrium.* ∎

Proof of Proposition 7

The conditions stated in Proposition 7 imply those in Proposition 6. Therefore it is immediate that $\{B_i\}_{i=1}^n$ is a *Beliefs Equilibrium.*

By construction of $B_i(\omega)$ and assumption $\triangle_i(\omega) = M(H_i(\omega)) \quad \forall \omega$, we have

$$\text{marg}_{S_j} B_i(\omega) = \text{convex hull of } \{f_j(\omega') \mid \omega' \in H_i(\omega)\} \qquad \forall \omega.$$

Now fix ω at the state where the rationality of the players and $\{u_i\}_{i=1}^n$ are *mutual knowledge* and $\{B_i\}_{i=1}^n$ is *common knowledge*. Player k *knows* player i's beliefs implies that $B_i(\omega') = B_i(\omega) \quad \forall \omega' \in$

$H_k(\omega)$ and therefore

$$\begin{aligned}
\text{convex hull of } \{f_j(\omega'') \mid \omega'' \in H_i(\omega')\} &= \text{marg}_{S_j} B_i(\omega') \\
&= \text{marg}_{S_j} B_i(\omega) \\
&= \text{convex hull of } \{f_j(\omega'') \mid \omega'' \in H_i(\omega)\} \quad \forall \omega' \in H_k(\omega).
\end{aligned}$$

Let $\Sigma_j^k \subseteq \{f_j(\omega') \mid \omega' \in H_k(\omega)\}$ be the set of extreme points of $\text{marg}_{S_j} B_k(\omega)$. I claim that $\Sigma_j^k \subseteq \text{marg}_{S_j} B_i(\omega)$ (and therefore $\text{marg}_{S_j} B_k(\omega) \subseteq \text{marg}_{S_j} B_i(\omega)$). Suppose that it were not true. Then there exists $\sigma_j \in \Sigma_j^k$ such that $\sigma_j \notin \text{marg}_{S_j} B_i(\omega') \quad \forall \omega' \in H_k(\omega)$. Therefore

$$\sigma_j \notin \cup_{\omega' \in H_k(\omega)} \{f_j(\omega'') \mid \omega'' \in H_i(\omega')\} \supseteq \{f_j(\omega') \mid \omega' \in H_k(\omega)\} \supseteq \Sigma_j^k \ni \sigma_j$$

which is a contradiction. Since i, j and k are arbitrary, we have $\text{marg}_{S_j} B_i(\omega) = \text{marg}_{S_j} B_k(\omega)$ and in particular, $\Sigma_j^i = \Sigma_j^k \equiv \Sigma_j$.

It only remains to show that

$$B_i(\omega) = \text{convex hull of} \{\sigma_{-i} \in M(S_{-i}) \mid \text{marg}_{S_j} \sigma_{-i} \in \Sigma_j\}.$$

$B_i(\omega)$ takes the form as required if and only if for each $\sigma_{-i} \in \times_{j \neq i} \Sigma_j$, there exists $\omega' \in H_i(\omega)$ such that $f_j(\omega') = \sigma_j \quad \forall j \neq i$.

Suppose that the condition stated above were not satisfied. Without loss of generality, assume that for player 1, there exists $\sigma_{-1} \in \times_{j \neq 1} \Sigma_j$ such that for each $\omega'' \in H_1(\omega)$, there exists $j \neq 1$ where $f_j(\omega'') \neq \sigma_j$. This implies that $\sigma_{-1} \notin B_1(\omega)$. $B_1(\omega)$ is *common knowledge* at ω implies that $B_1(\omega'') = B_1(\omega) \quad \forall \omega'' \in \mathcal{H}(\omega)$. Therefore $\sigma_{-1} \notin B_1(\omega'') \quad \forall \omega'' \in \mathcal{H}(\omega)$. Therefore for each $\omega'' \in \mathcal{H}(\omega)$, there exists $j \neq 1$ where $f_j(\omega'') \neq \sigma_j$.

Now consider player 2. The last sentence in the previous paragraph implies that for $\omega' \in \mathcal{H}(\omega)$ such that $f_2(\omega') = \sigma_2$, for each $\omega'' \in H_2(\omega')$, there exists $j \in \{3, \ldots, n\}$ such that $f_j(\omega'') \neq \sigma_j$. Therefore $\prod_{j=3}^n \sigma_j \notin \text{marg}_{\times_{j=3}^n S_j} B_2(\omega')$. Again $B_2(\omega)$ is *common knowledge* at $\omega \in \Omega$ implies $B_2(\omega'') = B_2(\omega) \quad \forall \omega'' \in \mathcal{H}(\omega)$. Therefore $\prod_{j=3}^n \sigma_j \notin \text{marg}_{\times_{j=3}^n S_j} B_2(\omega'') \quad \forall \omega'' \in \mathcal{H}(\omega)$ and we can conclude that for each $\omega'' \in \mathcal{H}(\omega)$, there exists $j \in \{3, \ldots, n\}$ such that $f_j(\omega'') \neq \sigma_j$.

Repeat the same argument for players $3, \ldots, n$ to conclude that for each $\omega'' \in \mathcal{H}(\omega)$, $f_n(\omega'') \neq \sigma_n$. This contradicts the fact that $\sigma_n \in \Sigma_n$. ∎

Proof of Proposition 8

By construction of $B_i(\omega)$ and assumption $\triangle_i(\omega) = M(H_i(\omega))$, it follows that $\text{marg}_{S_j} B_i(\omega) =$ convex hull of $\{f_j(\omega') \mid \omega' \in H_i(\omega)\}$. In particular, $f_j(\omega) \in \text{marg}_{S_j} B_i(\omega)$. At ω, the fact that player j is *rational* implies $f_j(\omega) \in BR_j(\omega)(B_j(\omega))$. Therefore $\text{marg}_{S_j} B_i(\omega) \cap BR_j(\omega)(B_j(\omega)) \neq \emptyset$.
∎

Proof of Proposition 9

Set $E_j = f_j(\omega)$. By construction of $B_i(\omega)$ and assumption $\triangle_i(\omega) = M(H_i(\omega))$, it follows that $B_i(\omega) = M(\{f_j(\omega') \mid \omega' \in H_i(\omega)\})$. In particular, there exists a probability measure in $B_i(\omega)$ which attaches probability one to E_j. The only proper subset of E_j is the empty set which is of course attached probability zero by every measure in $B_i(\omega)$. Therefore E_j satisfies condition 1 in Definition 9. At ω, the fact that player j is *rational* implies $f_j(\omega) \in BR_j(\omega)(B_j(\omega))$. Therefore condition 2 in Definition 9 is also satisfied. This completes the proof that $\{B_1, B_2\}$ is a *Nash Equilibrium Under Uncertainty*. ∎

REFERENCES

Anscombe, F. J. and R. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199-205.

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

Aumann, R. and A. Brandenberger (1991): "Epistemic Conditions for Nash Equilibrium," Manuscript, Harvard Business School.

Brandenberger, A. (1992): "Knowledge and Equilibrium in Games," *Journal of Economic Perspectives*, 6, 83-101.

Camerer, C., and M. Weber (1992): "Recent Developments in Modelling Preference: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty*, 5, 325-370.

Crawford, V. (1990): " Equilibrium Without Independence," *Journal of Economic Theory*, 50, 127-154.

Dekel, E., Z. Safra and U. Segal (1991): "Existence and Dynamic Consistency of Nash Equilibrium with Non-Expected Utility Preferences," *Journal of Economic Theory*, 55, 229-246.

Dow, J. and S. Werlang (1991): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," Manuscript, London Business School.

Dow, J. and S. Werlang (1994): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64, 305-324.

Dow, J. and S. Werlang (1992): "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica*, 60, 197-204.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643-669.

Epstein, L. G. and T. Wang (1994): "Intertemporal Asset Pricing under Knightian Uncertainty," *Econometrica*, 62, 283-322.

Fan, K. (1953): "Minimax Theorems," *Proceedings of the National Academy of Sciences*, 39, 42-47.

Fudenberg, D. and J. Tirole (1991): *Game Theory.* The MIT Press.

Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141-153.

Hurwicz, L. (1951): "Optimality Criteria for Decision Making Under Ignorance," Cowles Commission Discussion Paper.

Klibanoff, P. (1993): "Uncertainty, Decision, and Normal Form Games," *Journal of Economic Theory*, forthcoming.

Machina, M. (1989): "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, 27, 1622-1668.

Machina, M. and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 745-780.

Nash, J. (1951): "Non-cooperative Games," *Annuals of Mathematics*, 54, 286-95.

Pearce, D. (1984): "Rationalizable Strategic Behaviour and the Problem of Perfection," *Econometrica*, 52, 1029-50.

Raiffa, H. (1961): "Risk, Ambiguity and the Savage Axioms: Comment," *Quarterly Journal of Economics*, 75, 690-694.

Savage, L. (1954): *The Foundations of Statistics.* New York: John Wiley.

Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571-581.

Segal, U. (1990): "Two-stage Lotteries Without the Reduction Axiom," *Econometrica*, 58, 349-377.

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