

Extensive Form Games with Uncertainty Averse Players¹

Kin Chung Lo

Department of Economics, University of Toronto,
Toronto, Ontario, Canada M5S 1A1

July, 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 the Ellsberg Paradox and relevant experimental findings demonstrating that the beliefs of a decision maker may not be representable by a probability measure, this paper generalizes Nash Equilibrium in finite extensive form games to allow for preferences conforming to the multiple priors model developed in Gilboa and Schmeidler (1989). The implications of this generalization for strategy choices and welfare are studied. *Journal of Economic Literature* Classification Numbers: C72, D81.

¹I especially thank Professor Larry G. Epstein for pointing out this topic, and for providing supervision and encouragement. Remaining errors are my responsibility.

1. INTRODUCTION

The subjective expected utility model axiomatized by Savage (1954) has been the most popular model for studying decision making under uncertainty. In this model, the beliefs of a decision maker are represented by a probability measure. However, the descriptive validity of the model has been questioned since Ellsberg (1961) presented his famous example that people typically prefer to bet on drawing a red ball from an urn containing 50 red and black balls each, than from an urn containing 100 red and black balls in unknown proportions.² In fact, the pattern of preferences exhibited in Ellsberg's example is ruled out by any model of preferences in which underlying beliefs are represented by a probability measure. (Machina and Schmeidler (1992) call this property "probabilistic sophistication".)

Motivated by the Ellsberg Paradox, which demonstrates that there are situations where the information possessed by a decision maker about the states of nature is too "vague" or "ambiguous" to be representable by a probability measure, two important and closely related models have been developed.³ Schmeidler (1989) develops the Choquet expected utility model, in which the beliefs of a decision maker are represented by a capacity or non-additive probability measure. Gilboa and Schmeidler (1989) develop the multiple priors model, in which the beliefs of a decision maker are represented by a set of additive probability measures. In the multiple priors model, a decision maker is said to be uncertainty averse if the set of probability measures representing his beliefs is not a singleton.

Although the Ellsberg Paradox only involves a single decision maker facing an exogenously specified environment, it is natural to think that uncertainty aversion is also common in decision making problems where more than one person is involved. Since 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 line of research has already started. For normal form games of complete information, Dow and Werlang (1994), Klibanoff (1993) and Lo (1995a,b) generalize Nash Equilibrium; and Epstein (1995) generalizes rationalizability and *a posteriori* equilibrium. For normal form games of incomplete information, Epstein and Wang (1994) establish the general theoretical justification for the Harsanyi style formulation for non-Bayesian players. All the above papers either adopt the multiple priors model or consider a class of preferences which includes both

²Variations of the Ellsberg Paradox have been confirmed by many experimental studies. See Camerer and Weber (1992) for a survey.

³See Camerer and Weber (1992) for a survey of other models.

the Choquet expected utility and multiple priors models as special cases. However, serious study on extensive form games with uncertainty averse players has not yet been carried out.

This paper proposes a new equilibrium concept, Multiple Priors Nash Equilibrium, that generalizes Nash Equilibrium in extensive form games to accommodate preferences conforming to the multiple priors model. This research is important because many games of economic interest are extensive form games. The generalization creates a framework that enables us to study the effects of uncertainty aversion on strategic interaction in situations which are dynamic in nature.

It is well known that when players are expected utility maximizers, the definition of Nash Equilibrium in normal form games can be directly applied to (the normal form representation of the) extensive form games. Therefore, as far as Nash Equilibrium is concerned, a separate treatment for extensive form games is not needed. However, a separate treatment is required in the present setting of uncertainty averse players. The main reason is as follows. In an extensive form game, a strategy of a player is a specification of action taken by the player at every information set at which he is supposed to move. For an expected utility maximizing player, a strategy which maximizes his utility at the beginning of the game (that is, before anyone has made any move) will continue to be optimal for him when he arrives at an information set that does not contradict his initial beliefs on his opponents' strategy choices. This property does not hold when player's beliefs are represented by a set of probability measures. Therefore, it is required to ensure that the strategy chosen in equilibrium by an uncertainty averse player be optimal, not necessarily at the beginning of the game, but rather at every information set that he thinks he will possibly reach when he carries out the strategy. Unlike Nash Equilibrium, Multiple Priors Nash Equilibrium is an extensive form solution concept. That is, two extensive form games with the same normal form can have different sets of Multiple Priors Nash Equilibria. In the concluding section, I point out that dependence on the extensive form is natural when players are uncertainty averse. All other features of Nash Equilibrium are essentially preserved by the generalization. Therefore, a comparison between the two equilibrium concepts constitutes a *ceteris paribus* study of the effects of uncertainty aversion on extensive form games.

The paper is organized as follows. Section 2 contains a brief review of the multiple priors model, a discussion of how it is extended to the sequential choice setting and finally, a review of the definition of extensive form games. Section 3 defines Nash Equilibrium and its generalization. Section 4 makes use of the generalized equilibrium concept to illustrate how uncertainty aversion affects players' strategy choices and welfare. Some concluding remarks are offered in section 5.

2. PRELIMINARIES

2.1 Static Choice

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 with finite supports. Let (X, Σ_X) be the space of outcomes and (Ω, Σ_Ω) the space of uncertainty. For the purpose of this paper, assume that Ω is a finite set. Let \mathcal{F} be the set of all 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 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, I also use $p \in M(X)$ to denote the constant act that yields p in every state of the world, $x \in X$ the degenerate probability distribution on x and $\omega \in \Omega$ the event $\{\omega\} \in \Sigma_\Omega$. 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 weak preference ordering over acts. The relations of strict preference and indifference are denoted by \succ and \sim respectively.

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 $u : M(X) \rightarrow \mathbf{R}$ and a unique, nonempty, closed and convex set Δ of probability measures on Ω such that for all $f, f' \in \mathcal{F}$,

$$f \succeq f' \Leftrightarrow \min_{p \in \Delta} \int_{\Omega} u \circ f dp \geq \min_{p \in \Delta} \int_{\Omega} u \circ f' dp. \quad (1)$$

It is convenient, but in no way essential, to interpret Δ as “representing the beliefs underlying \succeq ”; I provide no formal justification for such an interpretation. 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 horse race/roulette wheel 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 $\alpha \in (0, 1)$.

There are three issues regarding the multiple priors model that will be relevant when the model is applied to games. The first 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 ,

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

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 Δ attaches zero probability to T .

The second concerns the notion of strict monotonicity. Given any preference ordering \succeq over acts, say that \succeq is *strictly monotonic* in an event T if for any two acts f and f' ,

$$f(\omega) \sim f'(\omega) \quad \forall \omega \notin T \quad \text{and} \quad f(\omega) \succ f'(\omega) \quad \forall \omega \in T \implies f \succ f'.$$

If \succeq is expected utility preferences, then it is strictly monotonic in T if and only if the decision maker attaches positive probability to T . Therefore expected utility preferences are strictly monotonic in all non- \succeq -null events. If \succeq is represented by the multiple priors model, then it is strictly monotonic in T if and only if every probability measure in Δ attaches positive probability to T . In this paper, I impose the requirement that preferences which are represented by the multiple priors model be strictly monotonic in all non- \succeq -null events. That is, every probability measure in Δ has the same support.

Finally, the notion of stochastic independence will also be relevant. Suppose the set of states Ω is a product space $\Omega^1 \times \dots \times \Omega^n$. 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: $p = \times_{i=1}^n \text{marg}_{\Omega^i} p$ where $\text{marg}_{\Omega^i} p$ is the marginal probability measure of p on Ω^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 Δ . Let $\text{marg}_{\Omega^i} \Delta$ be the set of marginal probability measures on Ω^i as one varies over all the probability measures in Δ . That is,

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

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

$$\Delta = \text{closed convex hull of } \{\times_{i=1}^n p^i \mid p^i \in \text{marg}_{\Omega^i} \Delta \quad \forall i\}. \quad (2)$$

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

2.2 Sequential Choice

The multiple priors model described in section 2.1 is a model of static or “one-shot” choice. It has to be extended if we want to use it to deal with sequential choice problems.

Think of two “times” $t = 0$ and $t = 1$ at which choices are made. There is complete resolution of uncertainty at time $t = 2$. Suppose that at time $t = 0$, the preference ordering \succeq over the set of acts \mathcal{F} mapping Ω to $M(X)$ is represented by the multiple priors model defined in (1). At time $t = 1$, suppose the decision maker learns that the true state is in a non- \succeq -null event $T \subset \Omega$. The relevant primitives now include: the set of states (T, Σ_T) , the set of outcomes X (unchanged) and the set of acts \mathcal{F}_T on T . Assume that acts in \mathcal{F}_T are ranked by a preference ordering \succeq_T which is represented by the following utility function: there exists a unique, nonempty, closed and convex set Δ_T of probability measures on T such that for all $f, f' \in \mathcal{F}_T$,

$$f \succeq_T f' \Leftrightarrow \min_{q \in \Delta_T} \int_T u \circ f dq \geq \min_{q \in \Delta_T} \int_T u \circ f' dq. \quad (3)$$

There remains the issue of the relationship between Δ and Δ_T . A natural procedure to revise the beliefs of the decision maker is to rule out some of the priors in Δ and then update the rest according to Bayes rule. Following Gilboa and Schmeidler (1993), an updating rule is characterized by a function R of the form $(\Delta, T) \mapsto R(\Delta, T)$ for every nonempty, closed and convex $\Delta \subseteq M(\Omega)$ and for every non- \succeq -null $T \in \Sigma_\Omega$ such that $R(\Delta, T) \subseteq \Delta$ is a nonempty, closed and convex set of measures with $p(T) > 0$ for all $p \in R(\Delta, T)$. The beliefs of the decision maker over T are then represented by the set of probability measures

$$\begin{aligned} \Delta_T \equiv \{q \in M(T) \mid \exists p \in R(\Delta, T) \\ \text{such that } q \text{ is updated from } p \text{ using Bayes rule}\}. \end{aligned}$$

An updating rule of particular interest is the *maximum likelihood updating rule*:

$$R(\Delta, T) = \{p \in \Delta \mid p(T) = \max_{\tilde{p} \in \Delta} \tilde{p}(T)\}.$$

Gilboa and Schmeidler (1993) provide an axiomatization of this updating rule. They show that if \succeq can be simultaneously represented by the Choquet expected utility and multiple priors models, the maximum likelihood updating rule coincides with the Dempster-Shafer updating rule (see Shafer (1976)). Moreover the updating rule is *commutative* in the sense that the results of this rule are independent of the order in which information is gathered (see their Theorem 3.3). In this paper, I assume that Δ_T is derived from Δ using the maximum likelihood updating rule. However, note that only Proposition 2 depends on what updating rule is adopted.

Unfortunately, when \succeq_T is updated from \succeq as above, **preferences** do not satisfy the following dynamic consistency requirement unless Δ is a singleton: for **all** $T \in \Sigma_\Omega$ and for **all** $f, f' \in \mathcal{F}_T$,

$g \in \mathcal{F}_{\Omega \setminus T}$,

$$f \succeq_T f' \Leftrightarrow \begin{bmatrix} f(\omega) & \text{if } \omega \in T \\ g(\omega) & \text{if } \omega \notin T \end{bmatrix} \succeq \begin{bmatrix} f'(\omega) & \text{if } \omega \in T \\ g(\omega) & \text{if } \omega \notin T \end{bmatrix}. \quad (4)$$

Suppose that at time $t = 0$, the act $\begin{bmatrix} f(\omega) & \text{if } \omega \in T \\ g(\omega) & \text{if } \omega \notin T \end{bmatrix}$ is chosen out of some feasible set. If (4) is not satisfied, choice made at $t = 0$ may not be respected at $t = 1$. However, note the following two remarks.

First, violation of dynamic consistency is not specific to this updating rule. Epstein and Le Breton (1993) show that there does not exist \succeq_T which is represented by (3) such that \succeq and \succeq_T satisfy (4).

Second, note that the above dynamic consistency condition on preferences is strong in the sense that (4) is required to hold for **all** events T and for **all** acts f, f' and g . When an uncertainty averse decision maker is confronted with a particular sequential decision problem where there are only **some** acts available for choice at $t = 0$ and $t = 1$ respectively and only **some** events will possibly be realized at $t = 1$, it is possible that there may exist an **act** which is “dynamically consistent” in the sense that it is optimal for him to implement at every decision point which he thinks that he will possibly reach as he carries out the act. Consider the game in Figure 1. (For all the game trees presented in this paper, the vector of numbers at each terminal history refers to the utility payoffs to the players (player 1 first, player 2 second, etc.) and the notation I_{ij} refers to the j th information set of player i .)

Insert Figure 1 here

At I_{11} , player 1 is uncertain about which strategy player 2 is going to use. The space of uncertainty Ω for player 1 can therefore be regarded as $\{L, M, R\}$. Suppose player 1 is uncertainty averse with beliefs represented by the set of probability measures

$$B_1 = \{p \in M(\{L, M, R\}) \mid p(L) \geq 0.05, p(M) \geq 0.8, p(R) \geq 0.05\}.$$

Given the beliefs of player 1, the utility of any strategy which involves playing k at I_{11} is equal to 30 while the utility of any strategy which involves playing r at I_{11} is less than 30. Therefore at I_{11} , it is optimal for him to use a strategy which involves playing k at I_{11} . Although the set of probability measures B_1 is not a singleton, such a strategy is “dynamically consistent”. After player 1 plays k at I_{11} , it will exclude the possibility of reaching the other two information sets I_{12}

and I_{13} . He simply does not have a chance to deviate from his original plan. See later sections for less extreme examples in which players actually move more than once.

2.3 Extensive Form Games

In this section, I define finite extensive form games of perfect recall. Formally, I need the following notation which is adapted from Osborne and Rubinstein (1994, p.200). An extensive form game has the following components:

- $\mathcal{N} = \{1, \dots, N\}$ is the set of players.
- H is the set of histories with typical element h . h is a sequence of actions taken by the players. The empty sequence is an element of H . If $h = (a^k)_{k=1, \dots, K} \in H$ and $K' < K$, then $(a^k)_{k=1, \dots, K'} \in H$. The latter is called a *subhistory* of the former. $(a^k)_{k=1, \dots, K}$ is *terminal* if there is no a^{K+1} such that $(a^k)_{k=1, \dots, K+1} \in H$. The set of all terminal histories is denoted Z . The set of actions available after a non-terminal history h is denoted $A(h) = \{a : (h, a) \in H\}$.
- \mathcal{P} is a function which assigns to each non-terminal history a member of \mathcal{N} . That is, $\mathcal{P}(h)$ is the player who is going to move after the history h .
- \mathcal{I}_i is a partition of $\{h \in H \setminus Z \mid \mathcal{P}(h) = i\}$ with typical element I_i . That is, \mathcal{I}_i is the collection of information sets for player i . It is required that for every $I_i \in \mathcal{I}_i$, $A(h) = A(h')$ for all $h, h' \in I_i$. Therefore we can define $A(I_i) \equiv A(h)$ for any $h \in I_i$. That is, $A(I_i)$ is the set of actions available to player i at his information set I_i .
- $u_i : Z \rightarrow \mathbf{R}$ is the payoff function of player i .

An extensive form game G is a tuple $\{\mathcal{N}, H, \mathcal{P}, \{\mathcal{I}_i\}_{i \in \mathcal{N}}, \{u_i\}_{i \in \mathcal{N}}\}$.⁴ It is assumed that G is common knowledge among the players.

Player i 's strategy space is $S_i \equiv \times_{I_i \in \mathcal{I}_i} A(I_i)$ with typical element s_i . That is, s_i is a function that assigns the action $s_i(I_i)$ to each information set $I_i \in \mathcal{I}_i$. Define $s_i(h) \equiv s_i(I_i) \quad \forall h \in I_i \quad \forall I_i \in \mathcal{I}_i$.

⁴To simplify notation, I assume throughout that the extensive form game does not involve any move by nature. For extensive form games involving moves by nature, \mathcal{P} is redefined to be a function which assigns to each non-terminal history a member of $\mathcal{N} \cup \{c\}$. If $\mathcal{P}(h) = c$, nature determines the action after the history h . f_c is a function which associates with every history h for which $\mathcal{P}(h) = c$ a probability measure $f_c(\cdot \mid h)$ on $A(h)$, where each such probability measure is independent of every other measure. We can treat nature as a player with constant payoff at every terminal history and assume that the beliefs of each player about nature's move coincide with f_c . All definitions and results below continue to be valid. Finally, see Osborne and Rubinstein (1994, p.203) for the definition of perfect recall.

Given any history $h = (a^1, \dots, a^l, a^{l+1}, \dots, a^K)$, s_i is *consistent* with h if for every subhistory (a^1, \dots, a^l) of h for which $\mathcal{P}(a^1, \dots, a^l) = i$, we have $s_i(a^1, \dots, a^l) = a^{l+1}$. s_i is *consistent* with an information set if there exists a history h in the information set such that s_i is consistent with h .

Throughout the paper, any statement concerning players i, j and k is intended for $i = 1, \dots, N$ and $i \neq j \neq k$.

3. EQUILIBRIUM CONCEPTS

3.1 Nash Equilibrium

The definition of Nash Equilibrium is well known; one version is stated in Definition 1 below. The main body of this section is to present Nash Equilibrium in a form that can be readily extended for our purposes. This is intended to convince the reader that the generalization undertaken in section 3.2 is appropriate.

Before anyone has made a move, player i is uncertain about the strategy choices of other players. Therefore $S_{-i} \equiv \times_{j \neq i} S_j$ can be regarded as the state space for player i . Each strategy $s_i \in S_i$ of player i can be regarded as an act over this state space. If player i plays s_i and the other players play $s_{-i} \in S_{-i}$, i receives the utility outcome $u_i(s_i, s_{-i}) \equiv u_i(z)$ where z is the unique terminal history such that every element in (s_i, s_{-i}) is consistent with z . According to the subjective expected utility model, player i 's beliefs over S_{-i} are represented by a probability measure b_i . Nash Equilibrium can be stated as follows:

Definition 1. $\{b_i\}_{i=1}^N$ is a *Nash Equilibrium* if the following conditions are satisfied.

1. *Agreement and Stochastic Independence.* There exists a probability measure $\beta(I_i) \in M(A(I_i))$ $\forall I_i \in \mathcal{I}_i \forall i = 1, \dots, N$ such that

$$b_i = \times_{j \neq i} \times_{I_j \in \mathcal{I}_j} \beta(I_j).$$

2. *Rationality.* Every $s_i \in \Gamma_i$ (the support of $\text{marg}_{S_i} b_j$) satisfies

$$\sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) b_i(s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} u_i(\hat{s}_i, s_{-i}) b_i(s_{-i}) \quad \forall \hat{s}_i \in S_i.$$

The interpretation of Definition 1 is as follows. Condition 1 is a restriction on players' initial beliefs. It can be broken into three parts:

- (i) $\text{marg}_{S_k} b_i = \text{marg}_{S_k} b_j$,
- (ii) $b_i = \times_{j \neq i} \text{marg}_{S_j} b_i$ and
- (iii) $\text{marg}_{S_j} b_i = \times_{I_j \in \mathcal{I}_j} \beta(I_j)$.

(i) says that the marginal beliefs of players i and j on the strategy choice of player k agree. (ii) says that the beliefs of player i about the strategy choices of all the other players are stochastically independent. (iii) says that the beliefs of player i about the moves of player j at all of j 's information sets are stochastically independent. The expression $\sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) b_i(s_{-i})$ in condition 2 is player i 's *ex ante* utility of the strategy s_i given his beliefs b_i . That is, it is player i 's utility of choosing the strategy s_i before any one has made any move. Therefore in a Nash Equilibrium $\{b_i\}_{i=1}^N$, every strategy s_i in Γ_i (which is the support of $\text{marg}_{S_i} b_j$) maximizes the *ex ante* utility of player i . That is, player j *knows* (in the sense defined in section 2.1) that player i will pick a strategy to maximize *ex ante* utility.

However, the most important property for $\{b_i\}_{i=1}^N$ to be an equilibrium, which is not explicit in Definition 1, is that every strategy $s_i \in \Gamma_i$ is optimal for player i to implement at every history that he thinks (according to his initial beliefs b_i) he will possibly reach when he carries out s_i . If this property were not satisfied, a strategy that is optimal and chosen by player i before anyone has made any move, may not be implemented by player i as the players proceed to play the game. The above suggests a reformulation of Nash Equilibrium in terms of the *interim* utility of players; that is, the utility of each player at points where he has to take an action.

Let me proceed to present Nash Equilibrium in terms of *interim* utilities. Define $\Gamma_{-i} \subseteq S_{-i}$ to be the support of b_i . Fix $s_i \in S_i$ and define

$$\mathcal{I}_i(s_i, \Gamma_{-i}) \equiv \{I_i \in \mathcal{I}_i \mid \exists h \in I_i \text{ and } \exists s_{-i} \in \Gamma_{-i} \text{ such that every strategy in } (s_i, s_{-i}) \text{ is consistent with } h\}.$$

That is, $I_i \in \mathcal{I}_i(s_i, \Gamma_{-i})$ if and only if there exists $s_{-i} \in \Gamma_{-i}$ such that player i will find himself at I_i when he implements s_i and his opponents implement s_{-i} . If player i has been carrying out the strategy s_i and finds himself at an information set $I_i \in \mathcal{I}_i(s_i, \Gamma_{-i})$, then at this point, he excludes all strategy profiles in Γ_{-i} that do not reach I_i . To be precise, player i retains only the strategy profiles of his opponents in

$$\Gamma_{-i}(I_i) \equiv \{s_{-i} \in \Gamma_{-i} \mid \exists h \in I_i \text{ such that every strategy in } s_{-i} \text{ is consistent with } h\}.$$

At I_i , we can regard $\Gamma_{-i}(I_i)$ as the state space for player i . The beliefs of player i over $\Gamma_{-i}(I_i)$ are represented by the probability measure b_{I_i} which is updated from b_i using Bayes rule. The objective of player i at this point is to choose a strategy \hat{s}_i (which may be different from s_i) that is consistent with I_i to maximize his *interim* utility

$$\sum_{s_{-i} \in \Gamma_{-i}(I_i)} u_i(\hat{s}_i, s_{-i}) b_{I_i}(s_{-i}).$$

I am now in a position to state the second definition of *Nash Equilibrium*.

Definition 2. $\{b_i\}_{i=1}^N$ is a *Nash Equilibrium* if the following conditions are satisfied.

1. *Agreement and Stochastic Independence.* There exists a probability measure $\beta(I_i) \in M(A(I_i))$ $\forall I_i \in \mathcal{I}_i \forall i = 1, \dots, N$ such that

$$b_i = \times_{j \neq i} \times_{I_j \in \mathcal{I}_j} \beta(I_j).$$

2. *Dynamic Consistency.* Every $s_i \in \Gamma_i$ (the support of $\text{marg}_{S_i} b_j$) satisfies

$$\sum_{s_{-i} \in \Gamma_{-i}(I_i)} u_i(s_i, s_{-i}) b_{I_i}(s_{-i}) \geq \sum_{s_{-i} \in \Gamma_{-i}(I_i)} u_i(\hat{s}_i, s_{-i}) b_{I_i}(s_{-i})$$

$\forall \hat{s}_i \text{ consistent with } I_i \quad \forall I_i \in \mathcal{I}_i(s_i, \Gamma_{-i}).$

Condition 1 in Definition 2 is the same as that in Definition 1. The interpretation of condition 2 is as follows. As the players proceed to play the game, every strategy $s_i \in \Gamma_i$ of player i must be a best response for player i (according to his *interim* utility) at every information set I_i that can be reached by a strategy profile in $s_i \times \Gamma_{-i}$. Call a strategy having this property *dynamically consistent*. Since Γ_i is the support of $\text{marg}_{S_i} b_j$ which in turn represents the marginal beliefs of player j on i 's strategy choice, player j knows that player i will choose a dynamically consistent strategy.

It is well known that Definitions 1 and 2 are equivalent due to the fact that expected utility preferences are dynamically consistent. Since Definition 1 is much easier to state, it is the more common formulation of Nash Equilibrium. As pointed out in section 2.2, preferences which are representable by the multiple priors model are not dynamically consistent. Therefore, the distinction between *ex ante* and *interim* utility maximization is important for the generalization pursued here.

3.2 Generalization of Nash Equilibrium

At the beginning of the game, player i is uncertain about the strategy choices of opponents. Since player i is uncertainty averse, beliefs over S_{-i} are represented by a closed and convex set of probability measures B_i . Assume that the probability measures in B_i have the same support $\Gamma_{-i} \subseteq S_{-i}$. This ensures that the preference ordering of player i is strictly monotonic in every non-null event. $\mathcal{I}_i(s_i, \Gamma_{-i})$ and $\Gamma_{-i}(I_i)$ are defined as in section 3.1. The beliefs of player i at $I_i \in \mathcal{I}_i(s_i, \Gamma_{-i})$ are represented by the set of probability measures B_{I_i} which is derived from B_i using the maximum likelihood updating rule. That is,

$$B_{I_i} \equiv \{q_i \in M(\Gamma_{-i}(I_i)) \mid \exists b_i \in B_i \text{ such that} \\ b_i(\Gamma_{-i}(I_i)) = \max_{\tilde{b}_i \in B_i} \tilde{b}_i(\Gamma_{-i}(I_i)) \text{ and } q_i \text{ is updated from } b_i \text{ using Bayes rule}\}. \quad (5)$$

The following is a generalization of Nash Equilibrium:

Definition 3. $\{B_i\}_{i=1}^N$ is a *Multiple Priors Nash Equilibrium* if the following conditions are satisfied.

1. *Agreement and Stochastic Independence.* There exists a closed and convex set of probability measures $\mathcal{B}(I_i) \subseteq M(A(I_i)) \forall I_i \in \mathcal{I}_i \forall i = 1, \dots, N$ such that

$$B_i = \text{closed convex hull of } \{\times_{j \neq i} \times_{I_j \in \mathcal{I}_j} \beta(I_j) \mid \beta(I_j) \in \mathcal{B}(I_j) \quad \forall I_j \in \mathcal{I}_j \quad \forall j \neq i\}.$$

2. *Dynamic Consistency.* Every $s_i \in \Gamma_i$ (the support of every probability measure in $\text{marg}_{S_i} B_j$) satisfies

$$\min_{b_{I_i} \in B_{I_i}} \sum_{s_{-i} \in \Gamma_{-i}(I_i)} u_i(s_i, s_{-i}) b_{I_i}(s_{-i}) \geq \min_{b_{I_i} \in B_{I_i}} \sum_{s_{-i} \in \Gamma_{-i}(I_i)} u_i(\hat{s}_i, s_{-i}) b_{I_i}(s_{-i}) \\ \forall \hat{s}_i \text{ consistent with } I_i \quad \forall I_i \in \mathcal{I}_i(s_i, \Gamma_{-i}).$$

The interpretation of Definition 3 parallels that of Definition 2. Condition 1 can again be broken into three parts:

(i) $\text{marg}_{S_k} B_i = \text{marg}_{S_k} B_j$,

(ii) $B_i = \times_{j \neq i} \text{marg}_{S_j} B_i$ and

(iii) $\text{marg}_{S_j} B_i = \text{closed convex hull of } \{\times_{I_j \in \mathcal{I}_j} \beta(I_j) \mid \beta(I_j) \in \mathcal{B}(I_j) \quad \forall I_j \in \mathcal{I}_j\}$.

(i) says that the marginal beliefs of the players agree. (ii) and (iii) require the beliefs of the players to be stochastically independent (in the sense defined in the last paragraph of section 2.1). Condition 2 in Definition 3 only differs from that in Definition 2 by allowing players' utility functions to be represented by the multiple priors model. Any Nash Equilibrium is also a Multiple Priors Nash Equilibrium. Therefore existence of the latter is ensured.

4. DOES UNCERTAINTY AVERSION MATTER?

4.1 Questions

In section 3, I have developed an equilibrium concept to study the effects of uncertainty aversion in the context of extensive form games. My objective here is to address the following two specific questions:

1. As an outside observer, one only observes the actions actually taken by the players, but not their beliefs. Is it possible for an outside observer to distinguish uncertainty averse players from Bayesian players?⁵
2. Does uncertainty aversion make the players worse off (better off)?

Let me clarify the nature of welfare comparison underlying question 2. Consider the following single person decision problem: let $\mathcal{G} \subseteq \mathcal{F}$ be a set of acts defined on a state space Ω . A decision maker is allowed to choose an act from \mathcal{G} . Suppose that initially, beliefs of the decision maker over Ω are represented by a probability measure \hat{p} and next that beliefs change from \hat{p} to the set of priors Δ . Given $f \in \mathcal{G}$, let $CE_{\Delta}(f)$ be the certainty equivalent of f , that is, $u(CE_{\Delta}(f)) = \min_{p \in \Delta} \int u \circ f dp$. Similar meaning is given to $CE_{\hat{p}}(f)$. Then uncertainty aversion makes the decision maker worse off (better off) if

$$\max_{f \in \mathcal{G}} CE_{\Delta}(f) \leq (\geq) \max_{f \in \mathcal{G}} CE_{\hat{p}}(f).$$

That is, fix a particular decision problem, \hat{p} and Δ . Uncertainty aversion makes the decision maker worse off (better off) if the certainty equivalent of participating in the decision problem is lower (higher) when beliefs are represented by Δ rather than \hat{p} . Roughly speaking, the same criterion applies to the context of extensive form games. Fix a Nash Equilibrium and a Multiple Priors Nash Equilibrium of a game. Uncertainty aversion makes player i worse off (better off) if the certainty equivalent of player i in the Multiple Priors Nash Equilibrium is (lower) higher than that in the

⁵In this paper, Bayesian means subjective expected utility maximizer.

Nash Equilibrium. More precise nature of the welfare comparison for different classes of games will be spelled out explicitly in sections 4.2 and 4.3.

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 Δ 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 $\omega \in \Omega$, $f \succeq g \Leftrightarrow f \succeq' g \Leftrightarrow p \succeq^* q$.

In section 4.2, I first examine the above questions in the context of normal form games (or extensive form games where players do not learn anything about the strategy choices of opponents when the game is played.) In section 4.3, I address the questions in the context of “proper” extensive form games. Various examples are constructed to demonstrate that uncertainty aversion leads to new predictions and enhances players’ welfare. On the other hand, I identify two specific classes of games for which the answers are opposite.

4.2 Normal Form Games

Say that G is a *normal form game* if for every terminal history z and every information set I_i in G , there exists a subhistory h of z such that $h \in I_i$. That is, when player i arrives at the information set I_i , he does not learn anything about the strategy choices of his opponents. When G is a normal form game, there is no updating of beliefs. Condition 2 in the definition of Multiple Priors Nash Equilibrium collapses to the following normal form restriction which parallels that in Definition 1:

- *Rationality.* Every $s_i \in \Gamma_i$ satisfies

$$\min_{b_i \in B_i} \sum_{s_{-i} \in \Gamma_{-i}} u_i(s_i, s_{-i}) b_i(s_{-i}) \geq \min_{b_i \in B_i} \sum_{s_{-i} \in \Gamma_{-i}} u_i(\hat{s}_i, s_{-i}) b_i(s_{-i}) \quad \forall \hat{s}_i \in S_i.$$

This normal form solution concept differs from those in Klibanoff (1993) and Lo (1995a). In particular, it is a strengthening of those in Dow and Werlang (1994) and Lo (1995b) by imposing the condition that the probability measures in B_i have the same support. Therefore I pause to examine its properties.⁶

⁶See Lo (1995a) for a detailed comparison of the solution concepts in Dow and Werlang (1994), Klibanoff (1993) and Lo (1995a).

Consider the following question in the context of single person decision making. As an outside observer, can we distinguish an uncertainty averse decision maker from a Bayesian decision maker? Under the following circumstance, the answer is “no”. Suppose that we observe an uncertainty averse decision maker who chooses an act f from a constraint set $\mathcal{G} = \{f, g\}$ that contains only two elements. 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 the state space $\Omega = \{\omega_1, \omega_2\}$. The feasible set of utility payoffs $\{(u(f(\omega_1)), u(f(\omega_2))), (u(g(\omega_1)), u(g(\omega_2)))\}$ generated by the constraint set $\{f, g\}$ are simply two points in \mathbf{R}^2 . To rationalize his choice by an expected utility function, we can draw a linear indifference curve which passes through $(u(f(\omega_1)), u(f(\omega_2)))$ and (lies above) $(u(g(\omega_1)), u(g(\omega_2)))$, with slope describing the probabilistically sophisticated beliefs of the decision maker. This argument immediately leads us to

Proposition 1. *Given a 2×2 normal form game, if $\{B_i, B_j\}$ is a Multiple Priors Nash Equilibrium, then there exist $b_i \in B_i$ and $b_j \in B_j$ such that $\{b_i, b_j\}$ is a Nash Equilibrium.*

Proposition 1 delivers two messages. The first regards the prediction of how the game will be played. Suppose $\{B_i, B_j\}$ is a Multiple Priors Nash Equilibrium of a 2×2 game. Recall that every probability measure in B_j has the same support Γ_i . The prediction associated with the equilibrium regarding strategies played is that i chooses some $s_i \in \Gamma_i$. According to Proposition 1, it is always possible to find at least one Nash Equilibrium $\{b_i, b_j\}$ such that the support of b_j is also Γ_i . Therefore the observed behavior of the uncertainty averse players (the actual strategies they choose) is also consistent with utility maximization given beliefs represented by $\{b_i, b_j\}$. This implies that an outsider who can only observe the actual strategy choice in the single game under study will not be able to distinguish uncertainty averse players from Bayesian players.

The second message regards the welfare of the players. The fact that the Nash Equilibrium is contained in the Multiple Priors Nash Equilibrium ($b_i \in B_i$) implies that

$$\max_{s_i \in S_i} \sum_{s_{-i} \in \Gamma_{-i}} u_i(s_i, s_{-i}) b_i(s_{-i}) \geq \max_{s_i \in S_i} \min_{\hat{b}_i \in B_i} \sum_{s_{-i} \in \Gamma_{-i}} u_i(s_i, s_{-i}) \hat{b}_i(s_{-i}).$$

The left hand side of the above inequality is the utility of player i in the Nash Equilibrium and the right hand side is that in the Multiple Priors Nash Equilibrium. In terms of certainty equivalent, the Nash Equilibrium Pareto dominates the Multiple Priors Nash Equilibrium. The game of matching pennies shows that the above inequality can be strict.

	L	R
U	10,0	0,10
D	0,10	10,0

Player i is the row player and player j is the column player. The unique Nash Equilibrium of this game is $\{b_i = (L, 0.5; R, 0.5), b_j = (U, 0.5; D, 0.5)\}$.⁷ The following is a Multiple Priors Nash Equilibrium:

$$B_i = \{p \in M(\{L, R\}) \mid 0.1 \leq p(L) \leq 0.9\} \text{ and } B_j = \{p \in M(\{U, D\}) \mid 0.1 \leq p(U) \leq 0.9\}.$$

The utility of the players in the Multiple Priors Nash Equilibrium is 1 and that in the Nash Equilibrium is 5.

The game below shows that Proposition 1 does not extend to two person normal form games where players have more than two strategies.⁸ It also shows that uncertainty aversion can make both players strictly better off.

	L	C	R
A	10,0	0,1	4,0
B	0,1	10,0	4,0
C	4,10	4,0	1 billion,4
D	4,0	4,10	1 billion,4

The unique Nash Equilibrium of this game is $\{b_i = (L, 0.5; C, 0.5), b_j = (A, 0.5; B, 0.5)\}$. In this Nash Equilibrium, the utility of player i is 5 and the utility of player j is 0.5. *Ex post*, player i receives at most 10 and player j at most 1. The following constitutes a Multiple Priors Nash Equilibrium:

$$B_i = R \text{ and } B_j = \{p \in M(S_i) \mid p(A) = 0, p(B) = 0, 0.3 \leq p(C) \leq 0.7\}.$$

In this Multiple Priors Nash Equilibrium, player i believes that player j will play R and therefore both C and D are i 's best responses. Player j does not know whether player i will play C or D . Since player j is uncertainty averse, he prefers to play R which ensures him the payoff of 4. As a result, player i receives 1 billion and player j receives 4 with certainty. Therefore the Multiple Priors Nash Equilibrium strongly Pareto dominates (both *ex ante* and *ex post*) the unique Nash Equilibrium of this game.

4.3 Extensive Form Games

The game in Figure 2 shows that uncertainty aversion in extensive form games can also lead to Pareto improvement.

⁷The notation $(L, 0.5; R, 0.5)$ refers to the probability measure which yields L with probability 0.5 and R with probability 0.5.

⁸Under the assumption that players have a strict incentive to randomize, Lo (1995a) proves that Proposition 1 holds for any two person normal form game in which players have any finite number of pure strategies.

Insert Figure 2 here

If player 2 is a Bayesian, his beliefs about what player 1 is going to do at the information set I_{12} are represented by a probability measure. The utility of the strategy k (m) is strictly higher than that of r if player 2 attaches probability of at least 0.5 that player 1 will take the action L (R) at the information set I_{12} . Therefore player 2 will never play r . This implies that if player 1 plays D , his payoff is equal to 0 with certainty. Therefore **any** solution concept which assumes that player 1 knows that player 2 is a Bayesian will predict that player 1 plays U and the payoff to each player is equal to 1 with certainty. The following constitutes a Multiple Priors Nash Equilibrium of the game:⁹

$$\mathcal{B}(I_{11}) = D, \quad \mathcal{B}(I_{12}) = \{\beta \in M(\{L, R\}) \mid 0.1 \leq \beta(L) \leq 0.9\}, \quad \mathcal{B}(I_{21}) = r.$$

In this equilibrium, player 2's beliefs about what player 1 is going to do at the information set I_{12} are represented by a set of probability measures $\mathcal{B}(I_{12})$. It predicts that player 1 plays D at I_{11} and player 2 plays r at I_{21} . Player 1 gets 1 billion with certainty and player 2 gets 1000 with certainty. Therefore, the Multiple Priors Nash Equilibrium strongly Pareto dominates (both *ex ante* and *ex post*) any Nash Equilibrium of this game. To see why uncertainty aversion leads to a better equilibrium, note that the probability measures in $\mathcal{B}(I_{12})$ that minimize player 2's utility of playing k and m are different. The probability measure that minimizes the utility of k is $(L, 0.1; R, 0.9)$ and the one that minimizes the utility of m is $(L, 0.9; R, 0.1)$. This makes playing k and m undesirable for player 2. On the other hand, if player 2 is a Bayesian, the two probability measures have to coincide. As a result, either k or m is strictly better than r .

Next consider the slightly modified game in Figure 3.

Insert Figure 3 here

The game tree is the same as that in Figure 2 except that each information set is a singleton. That is, it is a game of perfect information. The following constitutes a Multiple Priors Nash Equilibrium:

$$\begin{aligned} \mathcal{B}(I_{11}) &= D, \quad \mathcal{B}(I_{12}) = \{\beta \in M(\{L, R\}) \mid 0.1 \leq \beta(L) \leq 0.9\}, \\ \mathcal{B}(I_{13}) &= \{\beta \in M(\{L, R\}) \mid 0.1 \leq \beta(L) \leq 0.9\}, \quad \mathcal{B}(I_{21}) = r. \end{aligned}$$

⁹Whenever it is more convenient, I will state *Multiple Priors Nash Equilibrium* in terms of $\{\mathcal{B}(I_i)\}_{I_i \in \mathcal{I}, i=1, \dots, N}$. It is understood that $\{B_i\}_{i=1}^N$ is constructed according to condition 1 of Definition 3.

In this Multiple Priors Nash Equilibrium, player 2 is uncertainty averse about what player 1 is going to do at both the information sets I_{12} and I_{13} . This discourages player 2 to play either k or m . The equilibrium predicts that player 1 plays D at I_{11} and player 2 plays r at I_{21} . However, we can also construct a Nash Equilibrium to support the above prediction. For example,

$$\beta(I_{11}) = D, \quad \beta(I_{12}) = (L, 0.1; R, 0.9), \quad \beta(I_{13}) = (L, 0.9; R, 0.1), \quad \beta(I_{21}) = r$$

is such a Nash Equilibrium. Moreover, the two equilibria yield the same utility and therefore certainty equivalent to player 1 at I_{11} and player 2 at I_{21} .

In the Multiple Priors Nash Equilibrium constructed for the game in Figure 3, each player only moves once. The equilibrium constructed for the game in Figure 4 suggests that the above phenomenon of observational and welfare equivalence holds more generally.

Insert Figure 4 here

The following constitutes a Multiple Priors Nash Equilibrium of the game in Figure 4:

$$\mathcal{B}(I_{11}) = e, \quad \mathcal{B}(I_{12}) = q,$$

$$\mathcal{B}(I_{21}) = \{\beta \in M(\{g, h\}) \mid 0.1 \leq \beta(g) \leq 0.9\}, \quad \mathcal{B}(I_{22}) = \{\beta \in M(\{t, u\}) \mid 0.1 \leq \beta(t) \leq 0.9\}.$$

In this equilibrium, player 1 will take the action e at I_{11} , player 2 may take either g or h at I_{21} . If player 2 takes g , player 1 will take the action q at I_{12} . Therefore player 1 has the chance to move twice. Nevertheless, we can also construct a Nash Equilibrium that supports the above predictions and yields the same welfare to the players at every information set that is possibly reached. For example,

$$\beta(I_{11}) = e, \quad \beta(I_{12}) = q, \quad \beta(I_{21}) = (g, 0.1; h, 0.9), \quad \beta(I_{22}) = (t, 0.1; u, 0.9)$$

is such a Nash Equilibrium.

The messages conveyed by the games in Figures 3 and 4 can be clearly illustrated in the context of single person decision theory. First consider the following static choice scenario: suppose a decision maker is facing a state space $\Omega = \Omega^1 \times \dots \times \Omega^n$. His preference ordering \succeq over acts defined on Ω is represented by the multiple priors model defined in (1). Suppose that he is allowed to choose an act from $\{f^1, \dots, f^n\}$ where the payoff of f^i only depends on Ω^i for all $i = 1, \dots, n$. In what follows, f^i will also be identified in the obvious way as an act defined on Ω^i . In this case, \succeq

restricted to $\{f^1, \dots, f^n\}$ can be represented by the expected utility function $\int_{\Omega} u \circ f d(\hat{p}^1 \times \dots \times \hat{p}^n)$ where

$$\hat{p}^i \in \operatorname{argmin}_{p^i \in \operatorname{marg}_{\Omega^i} \Delta} \int_{\Omega^i} u \circ f^i dp^i \quad \forall i = 1, \dots, n. \quad (6)$$

Moreover, $\int_{\Omega} u \circ f^i d(\hat{p}^1 \times \dots \times \hat{p}^n) = \min_{p \in \Delta} \int_{\Omega} u \circ f^i dp$ for all $i = 1, \dots, n$.

The above argument immediately translates to the game in Figure 3 if we set the decision maker to be player 2 at I_{21} , $\Omega = \Omega^1 \times \Omega^2 \times \Omega^3$, $\Omega^1 = A(I_{12})$, $\Omega^2 = A(I_{13})$, Ω^3 a singleton, $f^1 = k$, $f^2 = m$, $f^3 = r$, $\hat{p}^1 = (L, 0.1; R, 0.9)$ and $\hat{p}^2 = (L, 0.9; R, 0.1)$.

Next, consider the following sequential choice scenario: think of two “times” at which choices are made. Suppose at time $t = 0$, the decision maker is facing the state space $\Omega = \Omega^0 \times \Omega^1 \times \dots \times \Omega^n$. Again, \succeq over acts defined on Ω is represented by the multiple priors model defined in (1) where the associated set of probability measures Δ takes the form of (2). Let $T = \omega^0 \times \Omega^1 \times \dots \times \Omega^n$ with $\omega^0 \in \Omega^0$ be a non- \succeq -null event. Recall that \succeq_T is represented by the multiple priors model defined in (3), where the associated set of probability measures Δ_T is derived from Δ using the maximum likelihood updating rule. Given an act f defined on Ω , $f_T \equiv [f(\omega) \text{ if } \omega \in T]$ denotes the induced act defined on T , and similarly for $f_{\Omega \setminus T}$. Let $\{f^1, \dots, f^n, c\}$ be a set of acts defined on Ω with the following properties:

1. The payoff of f_T^i only depends on Ω^i and that of $f_{\Omega \setminus T}^i$ only depends on $\Omega^0 \setminus \omega^0$ for all $i = 1, \dots, n$.
2. $f_{\Omega \setminus T}^i = f_{\Omega \setminus T}^j$ for all $i, j = 1, \dots, n$.
3. c is a constant act.

1 implies that f_T^i can be identified as an act defined on Ω^i and $f_{\Omega \setminus T}^i$ an act defined on $\Omega^0 \setminus \omega^0$. 2 implies that for all $p^0 \in M(\Omega^0)$ and for all $i, j = 1, \dots, n$, $\int_{\Omega^0 \setminus \omega^0} u \circ f_{\Omega \setminus T}^i dp^0 = \int_{\Omega^0 \setminus \omega^0} u \circ f_{\Omega \setminus T}^j dp^0$. In this particular setting, it is easy to verify that for all $i = 1, \dots, n$,

$$\begin{aligned} \min_{p \in \Delta} \int_{\Omega} u \circ f^i dp &= \min_{p^0 \in \operatorname{marg}_{\Omega^0} \Delta} [p^0(\omega^0) \min_{p^i \in \operatorname{marg}_{\Omega^i} \Delta} \int_{\Omega^i} u \circ f_T^i dp^i + \int_{\Omega^0 \setminus \omega^0} u \circ f_{\Omega \setminus T}^i dp^0] \\ &= \min_{p^0 \in \operatorname{marg}_{\Omega^0} \Delta} [p^0(\omega^0) \min_{q \in \Delta_T} \int_T u \circ f_T^i dq + \int_{\Omega^0 \setminus \omega^0} u \circ f_{\Omega \setminus T}^i dp^0]. \end{aligned}$$

The decision problem is specified as follows: at time $t = 0$, the decision maker has the option of choosing an act from $\{f^1, \dots, f^n, c\}$. If the decision maker chooses c , then the decision problem is

over. On the other hand, if the decision maker chooses any f^i , then at time $t = 1$, he will be told whether the true state is in T . If T does not happen, the decision problem is over. If T happens, he can switch to any act in $\{f^1, \dots, f^n\}$. The question is: suppose we observe that the decision maker chooses, for instance, f^2 at $t = 0$ and $t = 1$. Can this choice be rationalized as expected utility maximizing behavior? To see that the answer is “yes”, first go to $t = 1$. Observe that if T happens, the discussion of the static choice scenario in the previous paragraph is applicable. As in (6), pick $\hat{p}^i \in \operatorname{argmin}_{p^i \in \operatorname{marg}_{\Omega^i} \Delta} \int_{\Omega^i} u \circ f_T^i dp^i$ for all $i = 1, \dots, n$. Then go to $t = 0$ and pick

$$\hat{p}^0 \in \operatorname{argmin}_{p^0 \in \operatorname{marg}_{\Omega^0} \Delta} [p^0(\omega^0) \max_{i=1, \dots, n} \int_{\Omega^i} u \circ f_T^i d\hat{p}^i + \int_{\Omega^0 \setminus \omega^0} u \circ f_{\Omega \setminus T}^i dp^0]. \quad (7)$$

The expected utility functions $\int_{\Omega} u \circ f d(\hat{p}^0 \times \hat{p}^1 \times \dots \times \hat{p}^n)$ and $\int_T u \circ f_T d(\hat{p}^1 \times \dots \times \hat{p}^n)$ will support f^2 as an optimal choice at $t = 0$ and $t = 1$ respectively. Moreover, $\hat{p}^0 \times \dots \times \hat{p}^n \in \Delta$, $\int_{\Omega} u \circ f^2 d(\hat{p}^0 \times \hat{p}^1 \times \dots \times \hat{p}^n) = \min_{p \in \Delta} \int_{\Omega} u \circ f^2 dp$ and $\int_T u \circ f_T^2 d(\hat{p}^1 \times \dots \times \hat{p}^n) = \min_{q \in \Delta_T} \int_T u \circ f_T^2 dq$.

The relevance of this observation to the game in Figure 4 is as follows. Set the decision maker to be player 1, $\Omega = \Omega^0 \times \Omega^1 \times \Omega^2$, $\Omega^0 = A(I_{21})$, $\Omega^1 = A(I_{22})$, Ω^2 a singleton, $\omega^0 = g$, f^1 to be the strategy $\{s(I_{11}) = e, s(I_{12}) = r\}$, f^2 the strategy $\{s(I_{11}) = e, s(I_{12}) = q\}$ and c any strategy which involves taking the action d at I_{11} . At $t = 0$, player 1 is at I_{11} and facing the state space Ω . If he picks the strategy c , the game is over. On the other hand, if he picks f^1 or f^2 , then at $t = 1$, if the true state is in T , he will find himself at I_{12} . Given the Multiple Priors Nash Equilibrium, construct the Nash Equilibrium in the following manner: first go to $t = 1$. According to (6), $(t, 0.1; u, 0.9)$ is the probability measure in $\mathcal{B}(I_{22})$ which minimizes player 1’s utility of playing f_1 and rationalizes that f_2 is strictly better than f_1 . Then go to $t = 0$. According to (7), $(g, 0.1; h, 0.9)$ is the probability measure in $\mathcal{B}(I_{21})$ which minimizes player 1’s utility of playing f_2 .

The above intuition leads to

Proposition 2. *Given a two person game of perfect information, if $\{B_i, B_j\}$ is a Multiple Priors Nash Equilibrium, then there exist $b_i \in B_i$ and $b_j \in B_j$ such that $\{b_i, b_j\}$ is a Nash Equilibrium.*

Proof: See the appendix. ■

Same as Proposition 1, Proposition 2 delivers the message that in two person games of perfect information, uncertainty averse players are not observationally distinguishable from Bayesian players.

With regard to player’s welfare, the Nash Equilibrium constructed in the proof of Proposition 2 has the following property: the value function of player i (at every history h such that $\mathcal{P}(h) = i$

and there exists a strategy in Γ_j which is consistent with h) in the Nash Equilibrium is the same as that in the Multiple Priors Nash Equilibrium. This implies that for any Multiple Priors Nash Equilibrium, we can find a Nash Equilibrium that not only delivers the same predictions, but also yields the same *interim* utility and therefore certainty equivalent to each player.

Finally, the game in Figure 5 demonstrates that Proposition 2 does not extend to games of perfect information where more than two players are involved. This is due to the extra requirement imposed by the equilibrium concept that in a game with more than two players, the marginal beliefs of the players have to agree.

Insert Figure 5 here

The following constitutes a Multiple Priors Nash Equilibrium of the game:

$$\begin{aligned} \mathcal{B}(I_{11}) &= U, \quad \mathcal{B}(I_{12}) = \{\beta \in M(\{a, b\}) \mid 0.5 \leq \beta(b) \leq 0.8\}, \\ \mathcal{B}(I_{21}) &= (L, 0.5; R, 0.5), \quad \mathcal{B}(I_{31}) = (k, 0.5; r, 0.5). \end{aligned}$$

Note that in this Multiple Priors Nash Equilibrium, although the marginal beliefs of players 2 and 3 about what player 1 is going to do at the information set I_{12} are represented by the same set of probability measures $\mathcal{B}(I_{12})$, the probability measure in $\mathcal{B}(I_{12})$ that minimizes the utility of the strategy L for player 2 is $(a, 0.2; b, 0.8)$; the one that minimizes the utility of the strategy k for player 3 is $(a, 0.5; b, 0.5)$. In this Multiple Priors Nash Equilibrium, player 1 plays U at I_{11} and player 2 may play either R or L at I_{21} . However, the observation that player 1 plays U at I_{11} and player 2 plays R at I_{21} is incompatible with any Nash Equilibrium. To demonstrate this, note that player 1 will play U at I_{11} only if L is a best response for player 2 and r is a best response for player 3. R is a best response for player 2 only if k is a best response for player 3. Therefore we need to construct a Nash Equilibrium in which both k and r are best responses for player 3. This is true only if the probability that player 1 will play a at I_{12} is equal to 0.5. However it implies that R is not a best response for player 2.

5. CONCLUDING REMARKS

5.1 Refinements

This paper provides a generalization of Nash Equilibrium in extensive form games that allows preferences of players to be represented by the multiple priors model. One area of future research is

to provide refinements of the equilibrium concept. One refinement that can be readily formulated is subgame perfection: require (in the same way as for a Nash Equilibrium) the restriction of a Multiple Priors Nash Equilibrium $\{\mathcal{B}(I_i)\}_{I_i \in \mathcal{I}_i, i=1, \dots, N}$ of a game to any of its subgame to be a Multiple Priors Nash Equilibrium of that subgame. Note that all the Multiple Priors Nash Equilibria constructed in the examples of this paper are subgame perfect.

5.2 Normal vs. Extensive Form Solution Concepts

Unlike Nash Equilibrium, Multiple Priors Nash Equilibrium is an extensive form solution concept. Kohlberg and Mertens (1986) popularize the view that a solution concept should only depend on the reduced normal form of a game. An implicit assumption underlying their arguments is that players' preferences are represented by the expected utility model. Another area of research is to re-examine the validity of this view if we allow players' preferences to deviate from the expected utility model. For instance, consider the two games in Figures 6 and 7 which are "extensive form" versions of the Ellsberg Paradox.

Insert Figures 6 and 7 here

Suppose player 2 knows that player 1 draws a ball randomly from an urn which contains one red ball and two black and yellow balls in unknown proportions. Note that the two games have the same (reduced) normal form. Will player 2 play the two games in exactly the same way? In Figure 6, player 2 does not receive any information about the random draw. The Ellsberg Paradox suggests that if player 2 is uncertainty averse, he will strictly prefer to play L . In Figure 7, when player 2 is given the chance to move, he knows that the ball drawn out is not yellow. This piece of information may make him strictly prefer to play R . For instance, suppose that the initial beliefs of player 2 are represented by the set of probability measures $\{p \in M(\{R, B, Y\}) \mid p(R) = \frac{1}{3}, \frac{1}{6} \leq p(B) \leq \frac{3}{6}\}$. This predicts that player 2 will strictly prefer to play L in the game in Figure 6. Applying the maximum likelihood updating rule, his beliefs after receiving the information that the ball drawn out is not yellow are represented by the probability measure $(R, \frac{2}{5}; B, \frac{3}{5})$. This predicts that player 2 will strictly prefer to play R in the game in Figure 7. Therefore uncertainty aversion suggests the non-equivalence between the two extensive form games.

5.3 Choice of Strategy Space

Recall from section 2.1 that preferences represented by the multiple priors model are quasi-concave. This implies that the decision maker may have a strict incentive to randomize among

acts. However, the possibility of randomization was not considered in extensive form games with uncertainty averse players in section 3.2. It is therefore necessary to provide a clarification.

In the context of normal form games, Lo (1995a) clarifies the argument for and against the assumption that uncertainty averse players have a strict incentive to randomize. The argument goes 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 can be justified by the assumption that each player is dynamically consistent in the sense of Machina (1989) and perceives himself as moving last. On the other hand, we can understand the adoption of a pure strategy space as assuming that each player perceives himself as moving first.

In extensive form games, it is reasonable to assume that the perception of the players on the order of strategy choices agree with the order of moves which are explicitly specified by the game tree. For instance, if the game tree explicitly specifies that player 1 moves first and player 2 moves second. When player 1 decides what strategy to play, it is natural that player 1 perceives himself as the first person to make the strategy choice and therefore he will not have a strict incentive to randomize. This explains the choice of pure strategy space in this paper.

5.4 Behavioral Consistency

Finally, note that there is an alternative approach, termed *behavioral consistency*, that has been used to study extensive form games where players' preferences are not dynamically consistent (see, for example, Karni and Safra (1989)). This approach treats the same player i at each information set I_i as a distinct agent. The pure strategy space of agent I_i is therefore $A(I_i)$. Each agent I_i of the same player i has the same payoff at the terminal histories but he is only interested in choosing a strategy in $A(I_i)$ to maximize his own utility. Therefore the original game of N players is analyzed as a game of $\sum_{i=1}^N \#I_i$ agents. This approach also prevents a player from choosing a strategy that he knows will not be implemented. However it has the disadvantage that the equilibrium notion does not collapse to Nash Equilibrium when players are expected utility maximizers; see, for example, Myerson (1991, p.161). Since the purpose of my paper is to generalize Nash Equilibrium to accommodate uncertainty aversion, the behavioral consistency approach is not adopted here.

APPENDIX

Proof of Proposition 2. In a game of perfect information, every information set is a singleton. Therefore h will be used to denote both the history h and the information set that contains h .

Suppose $\{B_i, B_j\}$ is a Multiple Priors Nash Equilibrium. Recall that Γ_j denotes the support of every probability measure in B_i . For every history h such that $\mathcal{P}(h) = i$ and there exists a strategy in Γ_j which is consistent with h , define

$$V_i(h) \equiv \max_{s_i \in C_i(h)} \min_{b_h \in B_h} \sum_{s_j \in \Gamma_j(h)} u_i(s_i, s_j) b_h(s_j) \quad (8)$$

where $C_i(h) \subseteq S_i$ is the set of strategies of player i that is consistent with h . That is, $V_i(h)$ is the value function of player i at h . Suppose player i finds himself at h . The objective of player i is to pick a strategy $s_i \in C_i(h)$ such that the *interim* utility of s_i is equal to $V_i(h)$.

Since coalescing of moves does not affect the set of Multiple Priors Nash Equilibrium, there is no loss of generality by assuming that the extensive form game satisfies the following property: if $\mathcal{P}(h) = i$, then there does not exist an action a such that (h, a) is non-terminal and $\mathcal{P}((h, a)) = i$.

For every terminal history z , define $V_i(z) \equiv u_i(z)$. Suppose player i at h plays $s_i \in C_i(h)$ with $s_i(h) = \hat{a}$. There are three possibilities: (i) If (h, \hat{a}) is terminal, then player i receives $V_i(h, \hat{a})$. (ii) If $\mathcal{P}((h, \hat{a})) = j$ and player j plays $s_j \in \Gamma_j$ such that it is consistent with (h, \hat{a}) , $s_j((h, \hat{a})) = a$ and (h, \hat{a}, a) is terminal, then player i receives $V_i((h, \hat{a}, a))$. (iii) If $\mathcal{P}((h, \hat{a})) = j$ and player j plays $s_j \in \Gamma_j$ such that it is consistent with (h, \hat{a}) , $s_j((h, \hat{a})) = a$ and $\mathcal{P}((h, \hat{a}, a)) = i$, then player i will find himself at (h, \hat{a}, a) and he will have to pick a strategy in $C_i((h, \hat{a}, a))$ to maximize *interim* utility. According to condition 1 of Definition 3, the initial beliefs B_i of player i takes the form

$$B_i = \text{closed convex hull of } \{\times_{h' \in \mathcal{I}_j} \beta(h') \mid \beta(h') \in \mathcal{B}(h') \quad \forall h' \in \mathcal{I}_j\}.$$

According to the maximum likelihood updating rule defined in (5),

$$B_h = \text{closed convex hull of } \left. \begin{array}{l} \times_{h' \in \mathcal{I}_j} \beta(h') \mid \\ \beta(h') = a \quad \forall h' \in \mathcal{I}_j \text{ such that } (h', a) \text{ is a subhistory of } h \\ \beta(h') \in \mathcal{B}(h') \quad \forall h' \in \mathcal{I}_j \text{ such that } h' \text{ is not a subhistory of } h \end{array} \right\},$$

and similarly for $B_{(h, \hat{a}, a)}$. This implies that $\min_{b_h \in B_h} \sum_{s_j \in \Gamma_j(h)} u_i(s_i, s_j) b_h(s_j)$ can be rewritten as

$$\min_{\beta \in \mathcal{B}((h, \hat{a}))} \sum_{a \in A((h, \hat{a}))} \begin{cases} \beta(a) \min_{b_{(h, \hat{a}, a)} \in B_{(h, \hat{a}, a)}} \sum_{s_j \in \Gamma_j((h, \hat{a}, a))} u_i(s_i, s_j) b_{(h, \hat{a}, a)}(s_j) & \text{if } \mathcal{P}((h, \hat{a}, a)) = i \\ \beta(a) V_i((h, \hat{a}, a)) & \text{if } (h, \hat{a}, a) \text{ is terminal.} \end{cases}$$

(i), (ii) and (iii) imply that (8) can be rewritten as

$$V_i(h) = \max_{\hat{a} \in A(h)} \left\{ \begin{array}{l} \min_{\beta \in \mathcal{B}((h, \hat{a}))} \sum_{a \in A((h, \hat{a}))} V_i((h, \hat{a}, a)) \beta(a) \quad \text{if } \mathcal{P}((h, \hat{a})) = j \\ V_i((h, \hat{a})) \quad \text{if } (h, \hat{a}) \text{ is terminal} \end{array} \right\}. \quad (9)$$

It says that the objective of player i at h is to take an action \hat{a} in $A(h)$ such that the value of the expression inside the bracket in (9) is maximized. Take

$$\beta^{\min}((h, \hat{a})) \in \operatorname{argmin}_{\beta \in \mathcal{B}((h, \hat{a}))} \sum_{a \in A((h, \hat{a}))} V_i((h, \hat{a}, a)) \beta(a) \quad \text{if } \mathcal{P}((h, \hat{a})) = j.$$

Then we have

$$V_i(h) = \max_{\hat{a} \in A(h)} \left\{ \begin{array}{l} \sum_{a \in A((h, \hat{a}))} V_i((h, \hat{a}, a)) \beta^{\min}((h, \hat{a}))(a) \quad \text{if } \mathcal{P}((h, \hat{a})) = j \\ V_i((h, \hat{a})) \quad \text{if } (h, \hat{a}) \text{ is terminal} \end{array} \right\}. \quad (10)$$

Set $b_i = \times_{h' \in \mathcal{I}_j} \beta(h')$ where

$$\beta(h') = \begin{cases} \beta^{\min}(h') & \text{if there exists a strategy in } \Gamma_j \text{ which is consistent with } h' \\ \text{any element in } \mathcal{B}(h') & \text{otherwise.} \end{cases}$$

The equivalence of (9) and (10) implies that every $s_i \in \Gamma_i$ is optimal for player i to implement when player i 's initial beliefs are represented by the probability measure b_i . Since $b_j \in B_j$, the support of b_j is Γ_i . Therefore $\{b_i, b_j\}$ constitutes a Nash Equilibrium. ■

REFERENCES

- Anscombe, F. J. and R. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199-205.
- Camerer, C., and M. Weber (1992): "Recent Developments in Modelling Preference: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty*, 5, 325-370.
- Dow, J. and S. Werlang (1991): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64, 305-324.
- Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643-669.
- Epstein, L. G. (1995): "Preference, Rationalizability and Equilibrium," Manuscript, University of Toronto.
- Epstein, L. G. and M. Le Breton (1993): "Dynamically Consistent Beliefs must be Bayesian," *Journal of Economic Theory*, 61, 1-22.
- Epstein, L. G. and T. Wang (1994): "A Types Space for Games of Incomplete Information with Non-Bayesian Players," Manuscript, University of Toronto.
- Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141-153.
- Gilboa, I. and D. Schmeidler (1993): "Updating Ambiguous Beliefs," *Journal of Economic Theory*, 59, 33-49.
- Kohlberg, E. and J. F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003-1037.
- Karni, E. and Z. Safra (1989): "Ascending Bid Auctions with Behaviorally Consistent Bidders," *Annals of Operations Research*, 19, 435-446.
- Klibanoff, P. (1993): "Uncertainty, Decision, and Normal Form Games," *Journal of Economic Theory*, forthcoming.
- Kuhn, H. W. (1953): "Extensive Games and the Problem of Information," *Contributions to the Theory of Games*, II (Annals of Mathematics Studies), Princeton University Press.
- Lo, K. C. (1995a): "Equilibrium in Beliefs under Uncertainty," Manuscript, University of Toronto (first version: 1994).
- Lo, K. C. (1995b): "Nash Equilibrium without Mutual Knowledge of Rationality," Manuscript, University of Toronto.
- 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.
- Myerson, R. (1991): *Game Theory*, Harvard University Press.
- Osborne, M. J. and A. Rubinstein (1994): *A Course in Game Theory*, MIT Press.

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.

Shafer, G. (1976): *A Mathematical Theory of Evidence*. Princeton University Press.




Figure 1

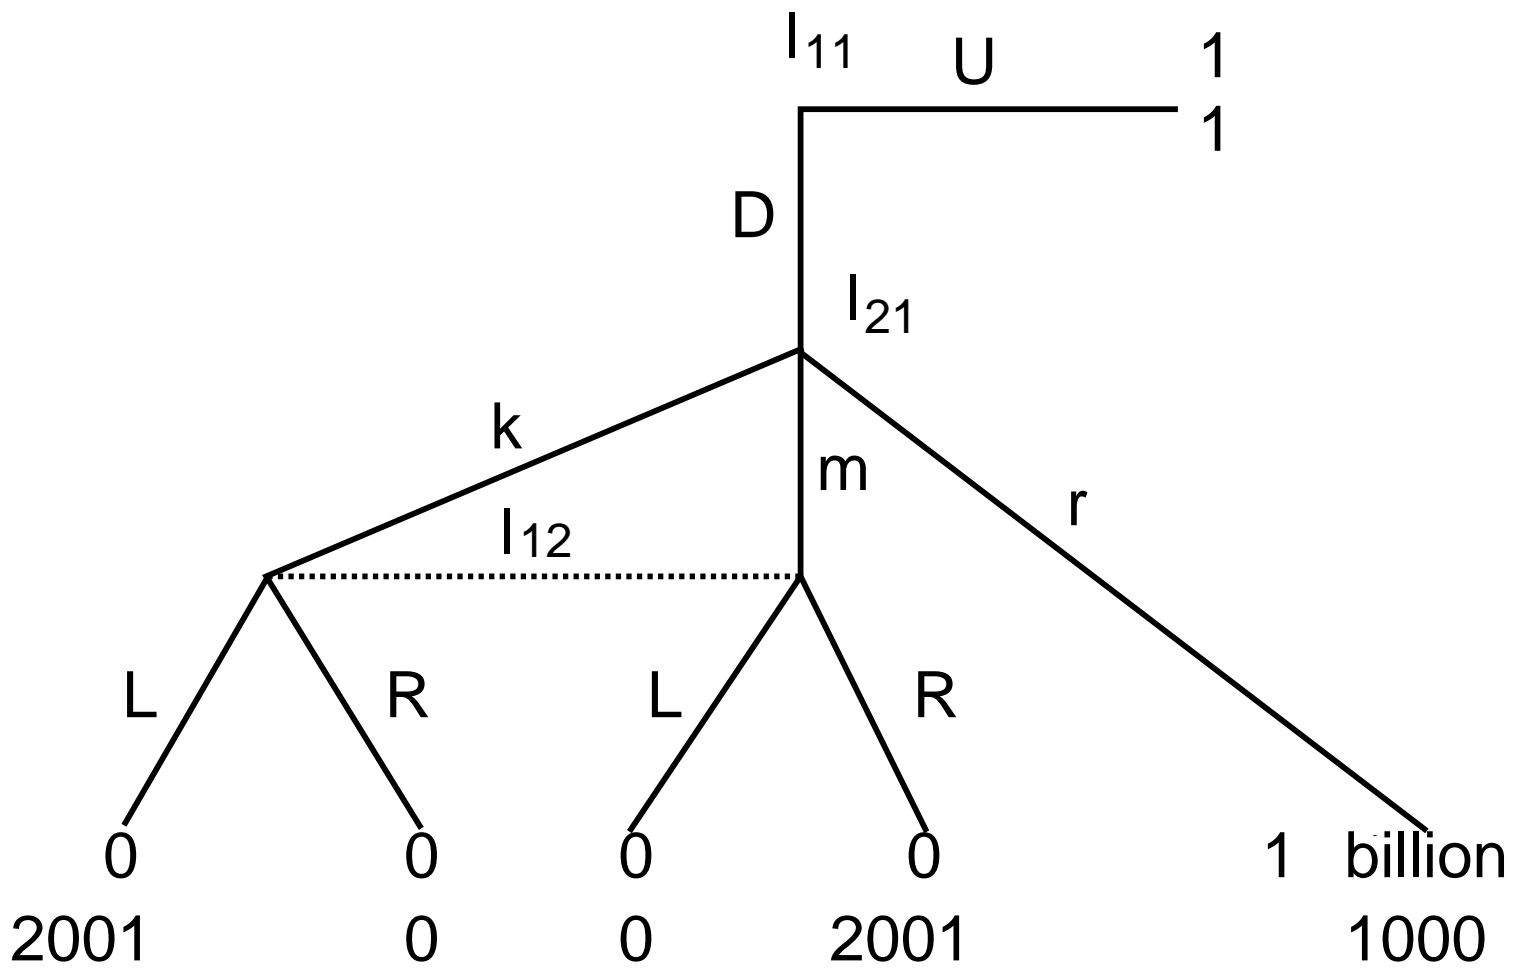


Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7