

Incomplete Information Games with Multiple Priors*

Atsushi Kajii
Institute of Economic Research
Kyoto University
kajii@kier.kyoto-u.ac.jp

Takashi Ui
Faculty of Economics
Yokohama National University
oui@ynu.ac.jp

First Draft: June 2003
This Version: November 2004

Abstract

We present a model of incomplete information games with *sets of priors*. Each player is endowed with a set of priors, and the sets of priors may or may not be common among the players. Upon arrival of private information, each player “updates” his set of priors to a set of posterior beliefs, by first applying the Bayes rule to each of priors in this set and then selecting a possibly proper subset of these. And finally the player evaluates his actions by the most pessimistic posterior beliefs à la Gilboa and Schmeidler (1989). So each player’s preferences may exhibit non-linearity in probabilities which can be interpreted as the player’s aversion to ambiguity or uncertainty. In this setup, we define a couple of equilibrium concepts, establish existence results for them, and demonstrate by examples how players’ views on uncertainty about the environment affect the strategic outcomes.

JEL classification: C72, D81, D82.

Key words: incomplete information games; multiple priors; ambiguity aversion; uncertainty aversion.

*An earlier version of this paper was circulated as KIER Working Paper 583, Kyoto University. We thank an anonymous referee for comments and the seminar participants at Osaka University, Otaru University of Commerce, Tokyo Institute of Technology, The Hong Kong University of Science & Technology, The University of Tokyo, Stanford University, The University of Paris I, The University of Venice, and The University of Mannheim. Kajii acknowledges financial support by MEXT, Grant-in-Aid for 21st Century COE Program. Ui acknowledges financial support by MEXT, Grant-in-Aid for Scientific Research.

1 Introduction

We present a model of incomplete information games with multiple priors. More specifically, our model is the same as the standard Bayesian games of incomplete information except for one point: instead of a single common prior over the states, we assume that there is a *set of priors* for each player which may differ among the players.

Upon arrival of private information, each player “updates” the set of priors to a set of posterior beliefs, by first applying the Bayes rule to each of priors in this set and then selecting a possibly proper subset of these. And finally the player evaluates his actions by the most pessimistic belief à la Gilboa and Schmeidler (1989). So each player’s preferences may exhibit non-linearity in probabilities which can be interpreted as the player’s aversion to uncertainty or ambiguity.

Our model is simple and very tractable since it is a minimal departure from the standard Bayesian approach for the games of incomplete information. On the other hand, our model is rich enough to capture important aspects of incomplete information games, and distinguish the two key ingredients of incomplete information: lack of information about payoffs and uncertainty (ambiguity) about them.

In this setup, we define a couple of equilibrium concepts, establish existence results for them, and demonstrate by examples how players’ views on uncertainty about the environment affect the strategic outcomes. We will make it clear that not only the aversion to uncertainty matters in one player’s own decision making, but also it has strategic effects on the other players’ strategic decision making.

To appreciate our contribution, let us start with a brief review of the standard approach. By definition, a game of incomplete information is a description of strategic environment where the players do not necessarily know some of the important parameters of the environment, such as payoffs. Harsanyi (1967–68) advocated the approach of representing incomplete information games by Bayesian games with imperfect information, which is now the standard approach.

For our purpose, Harsanyi’s points can be summarized into the following two points. The first point is that the source of uncertainty can be expressed by an underlying state space, and the incompleteness can be reduced to differences of private information among the players. The reason is that if there is any uncertainty about the fundamental specification of the game, it is due to the fact that the description of the underlying states is insufficient. The rational players then would not and should not be satisfied with such a description. If the description is refined, to its limit, then all the payoff relevant issues and the structure of knowledge among players can be summarized in a state space. We accept this view in our model, and assumes that there is an underlying state space and each state is a complete description of the game.

The second point is that there is a *single, common* prior over these states and each player evaluates his private information by the Bayes rule. We have a different view on this. When the players do not know payoffs and thus the information is incomplete, the players must take into account “risks” about payoffs: since his private information is insufficient, the player can learn payoffs only probabilistically at best. But there is also uncertainty

or ambiguity about the strategic environment: the player may have some thoughts about possibilities, but they are so vague that the player is unable to assign probabilities.

In principle, a model of incomplete information games should be able to address risks and uncertainty separately, to study if and how these two aspects of incomplete information affect the outcome of the game. But in Harsanyi's framework with a single and common prior, by construction, there is no technical difference between genuine lack of information and uncertainty about payoffs and/or knowledge among players.

So even accepting the postulate that the players should be able to describe all the relevant states, we think that it is still natural that the players have little idea about the likelihood of these states. For instance we want to allow for an event whose probability is perfectly agreed, but the players cannot evaluate subevents of the event; that is, there are many ways to assign probabilities to the subevents, and the players are not sure about which is the right one. Also, if such possibilities are to be considered, it is no longer clear if these assignments should be *common*. Thus we allow for *multiple* priors which are not necessarily common.

The assumption of multiple priors can be justified at the level of individual decision making to begin with (Gilboa and Schmeidler, 1989). It has been criticized that the standard framework of Bayesian theory does not necessarily capture the aspect of decision making which can be attributed to uncertainty or ambiguity of the problem; e.g., the Ellsberg paradox. Uncertainty is not resolved by considering more sophisticated Bayesian models, and the attitude towards uncertainty should be modelled differently from that towards risks.

In strategic environments of incomplete information, we think that the use of multiple priors is justified and interesting all the more. Say we accept the Bayesian view and assume that the Ellsberg paradox type problem does not occur as far as a player evaluates his own payoffs. But there still remains a room for uncertainty about what the other players might be contemplating, and a player may hesitate to assume the same ability for all the other players. The strategic decision making will be affected by the way the game is perceived by, possibly a fraction of, the other players in the game.

In other words, even if the uncertainty concerning payoff relevant issues may be reduced to risks, the same procedure may be unduly demanding for the uncertainty on knowledge about how the other players might think about each other. In such an environment, there will be strategic effects from the uncertainty about players' knowledge, which cannot be addressed by a single common prior model.

Allowing uncommon sets of priors is perhaps more controversial, since if these sets are singletons, our model is a Bayesian game with uncommon prior, and thus subject to Harsanyi's criticism. But since we have no firm foundation of the type space construction for our setup, we do not know if a common set of priors is a natural assumption in this context. So we chose to present a model allowing for personal sets of priors and established the existence results. But all of our motivating examples have a common set of priors. In other words, the strategic effects we presented in this paper do not rely on the diversity of sets of priors.

We shall now explain the other ingredients of our model analysis, i.e., updating rules. For this, it will be helpful to consider the following simple example. Say there are two experienced investors who are interested in the rating of the profitability of a firm to be announced near future. There are four possible ratings, 1, 2, 3, and 4, where the smaller number is better. From the past experiences, the investors know each other very well and they share the idea that the profitability of this firm is neutral in the sense that the probability of good ratings 1 or 2 is 0.5, and that of bad ratings 3 or 4 is 0.5. Investor 1 is very good at identifying very promising firm and Investor 2 is very good at identifying very bad one. That is, by the private information, Investor 1 can tell $\{1\}$ from $\{2, 3, 4\}$, and Investor 2 can tell $\{4\}$ from $\{1, 2, 3\}$. Now suppose that the firm's profitability is in fact very good. Then Investor 1 knows for sure that the rating will be 1, but he is unsure about what Investor 2 might think. Investor 1 can deduce that Investor 2 concludes that the firm's profitability is not very bad, but it is ambiguous how Investor 2 might think about the relative likelihood in $\{1, 2, 3\}$.

A natural set of priors is the set of all probabilities on $\{1, 2, 3, 4\}$ such that probabilities of events $\{1, 2\}$ and $\{3, 4\}$ are 0.5, and thereby we can express the fact that the two investors agree upon the likelihood of these two events. When the true state is 1, Investor 2 updates his beliefs. Since he was unable to determine the probabilities over 3 and 4 to start with, he will not be able to tell whether 4 did not occur by chance, or 4 had no chance to start with. So learning that 4 has not occurred does not reduce uncertainty about the relative probabilities of $\{1, 2\}$ and $\{3\}$, not to mention those between 1 and 2. Then Investor 1 must take into account not only the fact that Investor 2 does not know if 1 has occurred or not, but also the fact that Investor 2 is not able to assign a unique set of probabilities to $\{1, 2, 3\}$.

Notice that the above reasoning depends on the updating rule. In principle player's optimal behavior depends on how he updates priors. We assume that the updating rule is also part of the description of the game, i.e., the updating rules are exogenously given and the players understand the updating rules of all the players. We do so by the following reasons. First, in the multiple priors models, there are various reasonable updating rules, and they have different strategic implications. So the model should not subscribe itself to a particular updating rule. Secondly, since updating rules have strategic implications, endogenizing them will yield excessive degrees of freedom, and consequently the model will lose its descriptive power. Thirdly, studying the roles of different updating rules in games is of interest for itself.

Obviously, there can be many reasonable criteria for decision making under uncertainty or ambiguity. We adopt the Gilboa-Schmeidler approach to capture uncertainty and aversion to it. Although we are aware of valid criticism of Epstein and Zhang (2001) that this is not exactly what should be termed uncertainty or ambiguity, we contend that the Gilboa-Schmeidler approach at least addresses some important aspects of uncertainty and ambiguity, and its tractable form is also very appealing.

We propose interim equilibrium concepts in which each player chooses the best action for any realization of private signal in equilibrium. An ex ante equilibrium is defined

naturally from the priors as well in our model, but it is well known that a player with multiple priors tends to exhibit dynamic inconsistent behavior for any updating rule; that is, a strategy which is utility maximizing ex ante may specify actions which will be deemed inferior once private information is received, and vice versa. We contend that the interim notion is more relevant in applications; even in Harsanyi's framework, the ex ante maximization of utility is a purely theoretical tool and it happens to coincide with the interim notion because the expected utility model with Bayesian updating is dynamically consistent.

Nevertheless, it is of interest to compare the interim equilibrium with the ex ante equilibrium to investigate the implication of some players' dynamic inconsistent behavior, since one player's dynamically inconsistent behavior affects all the others who may well be textbook Bayesian players with single prior. We emphasize again that the study of such a non-Bayesian behavior is all the more important in strategic environments. Even if one regards the non-Bayesian behavior as irrelevant at the level of individual decision making, he will choose different actions from the ones he would against Bayesians, since actions taken by those non-Bayesians will influence his welfare, and vice versa.

Let us mention related works. Lo (1998) uses a multiple prior model to study auctions as a game of incomplete information, which Ozdenoren (2002) and Chen *et al.* (2003) further elaborate. Salo and Weber (1995) use a Choquet expected utility model with a convex capacity to study auctions as a game of incomplete information. Their model can be regarded as a game of incomplete information with multiple priors because it is known that a Choquet expected utility model with a convex capacity can be represented as a multiple prior model. As far as we know, however, there has been no attempt to study incomplete information games as a general class of games with multiple priors, with an important exception of Epstein and Wang (1996).¹ They present a very general form of games of incomplete information and our model constitutes a subclass of their games. Our focus is rather on presentation of workable and user friendly special models, which is rich enough to address issues of incompleteness of information beyond the standard Bayesian approach.

Let us conclude this introduction with an outline of this paper. We shall give the details of our model in Section 2. We then propose two equilibrium concepts, which are natural extensions of the Bayesian Nash equilibrium in Section 3. We contend that both concepts make sense, and establish the existence results for them. Section 4 contains examples to demonstrate how our model works and to show some interesting features of the model. In Section 5 we discuss the foundations of the model to argue why our model with these equilibrium concepts constitutes a desirable representation of incomplete information games. We also provide a result generalizing the agreement theorem of Aumann (1976) under the common multiple prior assumption, which is an interesting by-product. We also discuss complete information games and related works in the literature.

¹See also Ahn (2003).

2 Incomplete information games with multiple priors

2.1 Basic setup

We consider finite player incomplete information games with finitely many actions and states. Except for uncommon multiple priors, the setup is standard. The players are indexed by $i \in \mathcal{I} := \{1, \dots, I\}$. Each player $i \in \mathcal{I}$ has a finite set of actions denoted by A_i . The set of action profiles is denoted by $A = \prod_{i \in \mathcal{I}} A_i$ with generic element $a = (a_i)_{i \in \mathcal{I}}$. We shall also write $a = (a_i, a_{-i}) \in A_i \times \prod_{j \neq i} A_j$ abusing notation.

The set of payoff relevant states is denoted by Ω , and Ω is assumed to be finite.² The incompleteness of information is summarized by a random signal $\tau = (\tau_i)_{i \in \mathcal{I}}$, each component of which is observed privately by each player. When $\omega \in \Omega$ occurs, player $i \in \mathcal{I}$ observes a signal $\tau_i(\omega)$, and then chooses an action in A_i . Denote by T_i the range of τ_i and let $T = \prod_{i \in \mathcal{I}} T_i$. Thus τ_i is a function from Ω to T_i and $\tau = (\tau_i)_{i \in \mathcal{I}}$ is a function from Ω to T .

For any finite set X denote by $\Delta(X)$ the set of all probability distributions on X . A strategy of player $i \in \mathcal{I}$ is a function σ_i from T_i to $\Delta(A_i)$. Write $\sigma_i(a_i|t_i)$ for the probability of player i choosing action $a_i \in A_i$ when he observes $t_i \in T_i$ by convention. Denote by S_i the set of all strategies for player i and let $S = \prod_{i \in \mathcal{I}} S_i$ be the set of strategy profiles. For an action profile $a = (a_i)_{i \in \mathcal{I}} \in A$ and a profile of realization of signals $t = (t_i)_{i \in \mathcal{I}} \in T$, we write $\sigma(a|t)$ for the probability of action profile a chosen, i.e., $\sigma(a|t) = \prod_{i \in \mathcal{I}} \sigma_i(a_i|t_i)$. We shall also write $\sigma = (\sigma_i, \sigma_{-i}) \in S_i \times \prod_{j \neq i} S_j$ and $\sigma_{-i}(a_{-i}|t_{-i}) = \prod_{j \neq i} \sigma_j(a_j|t_j)$, abusing notation.

Player i 's preference ordering over strategy profiles will be generated by a payoff function $u_i : A \times \Omega \rightarrow \mathbb{R}$. In the standard incomplete information game, one could assume in addition that the payoff function $u_i(a, \omega)$ depends on players' observed realizations of the signals only, by replacing $u_i(a, \omega)$ with $\hat{u}_i(a, t) = E[u_i(a, \omega) | \tau(\omega) = t]$. In our framework, however, this transformation may change the strategic structure of the game because we will consider multiple priors and the expectation operator is not uniquely determined.

2.2 Multiple priors

We depart from the standard framework of incomplete information games by assuming that there is a non-empty compact set of priors $\mathcal{P}_i \subseteq \Delta(\Omega)$ for player $i \in \mathcal{I}$. We assume that there is no null signal, i.e., $P(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$, $P \in \mathcal{P}_i$, and $i \in \mathcal{I}$. The standard incomplete information game corresponds to the case where each \mathcal{P}_i is a singleton for all $i \in \mathcal{I}$, and they coincide. The set \mathcal{P}_i is intended to capture the uncertainty about the structure of the game, which is different from the strategic risk generated by the other players' choices of actions. We will demonstrate the different roles of ‘‘uncertainty’’ and ‘‘risk’’ by some examples later.

²The model can be extended to the model where Ω is an infinite measurable state space, but we restrict our attention to the finite case in order to avoid various measurability and continuity issues associated with an infinite state space.

A natural and interesting case is when the set of priors is generated by an underlying information sub-field $\mathcal{E} \subseteq 2^\Omega$ and a probability measure Q defined over \mathcal{E} . This case, which covers all of our illustrating examples presented later, naturally induces a common set of multiple priors. First, note that Q assigns a probability to every $E \in \mathcal{E}$, but not to $E \notin \mathcal{E}$. Thus, if $\mathcal{E} \neq 2^\Omega$, then a probability of some event is not known. The inner measure $P_* : 2^\Omega \rightarrow [0, 1]$ and the outer measure $P^* : 2^\Omega \rightarrow [0, 1]$ are defined by the rules:

$$P_*(E) = \sup_{X \subseteq E, X \in \mathcal{E}} Q(X), \quad P^*(E) = \inf_{E \subseteq X, X \in \mathcal{E}} Q(X)$$

for all $E \subseteq \Omega$. If $\mathcal{E} = 2^\Omega$, $P_* = Q = P^*$ by construction. Each player i has the set of priors as follows:

$$\mathcal{P}_i = \{P \in \Delta(\Omega) : P_*(E) \leq P(E) \leq P^*(E) \text{ for all } E \subseteq \Omega\}. \quad (1)$$

To interpret, think of \mathcal{E} as an information structure which is known to player i , and Q is a probability assessment of \mathcal{E} , and assume that they are objectively given. For an ‘‘unknown’’ event $E \notin \mathcal{E}$, $P_*(E)$ is the most cautious estimate of the probability of E and $P^*(E)$ is the most optimistic estimate of probability of E . Thus in this case, the set \mathcal{P}_i can be thought as the set of priors which are consistent with Q and \mathcal{E} in the sense that each $P \in \mathcal{P}_i$ assigns to each unknown event a probability weight at least as much as the cautious estimate and at most as much as the optimistic estimate.

2.3 Updating rules

Each player chooses an action after the private signal is revealed, as we mentioned earlier. Like in the standard Bayesian games, updating upon private information generates the differences in views of players. Since the prior is not unique, however, the private information will matter through two channels in our framework.

The first is the channel through the standard Bayesian updating: when player $i \in \mathcal{I}$ observes t_i , he calculates the conditional probability distribution for *each* prior $P \in \mathcal{P}_i$ by the Bayes rule, and the resulting distribution in turn affects his choice of actions, just as in the standard framework.

But the following second channel is not captured in the standard framework: since information is private, the sets of conditional probability distributions are different among the players in general, even if \mathcal{P}_i 's are common among $i \in \mathcal{I}$. In heuristic words, a revelation of private information might change the degree of uncertainty about the structure of the game, and this may occur differently among the players even with a common set of priors. Thus differences in views about the uncertainty of payoffs are generated by how the players process their private information to re-evaluate the uncertainty. We shall formalize these ideas below.

For each $P \in \Delta(\Omega)$ and $t_i \in T_i$, denote by $P(\cdot|t_i) \in \Delta(\Omega)$ the conditional probabilities over Ω ; that is, $P(E|t_i) = P(\tau_i^{-1}(t_i) \cap E) / P(\tau_i^{-1}(t_i))$ for $E \subseteq \Omega$. Let

$$\mathcal{P}_i(t_i) = \{P(\cdot|t_i) \in \Delta(\Omega) : P \in \mathcal{P}_i\}$$

be the set of conditional probability distributions when $t_i \in T_i$ has been observed. An *updating rule* $\Phi_i : T_i \rightarrow 2^{\mathcal{P}_i(t_i)}$ for player $i \in \mathcal{I}$ is a function that assigns a non-empty compact subset of $\mathcal{P}_i(t_i)$ to each $t_i \in T_i$. After t_i is observed, player i uses posteriors in $\Phi_i(t_i)$ to evaluate his actions. The updating rules for players are given as one of the primitives of the game, and the equilibrium concepts for themselves will be well defined for any such given rules.

When \mathcal{P}_i is a singleton for all $i \in \mathcal{I}$, the updating rule coincides with the Bayes rule, and our model will be reduced to the standard Bayesian games with possibly uncommon priors. But when \mathcal{P}_i is not a singleton, there is a vast variety of sensible updating rules in principle. Among them, there are a couple of natural and technically tractable updating rules of particular interest, and we shall concentrate on these cases in the examples we consider.³

The first is the Fagin-Halpern updating rule or the *full Bayesian* (FB) updating rule (Fagin and Halpern, 1990; Jaffray, 1992):

$$\Phi_i(t_i) = \mathcal{P}_i(t_i) \text{ for all } t_i \in T_i. \quad (2)$$

In words, this is the case where the private information leads the player to update the risk component of the priors but does not help him to update the degree of uncertainty.

The second is the Dempster-Shafer updating rule or the *maximum likelihood* (ML) updating rule (Dempster, 1967; Shafer, 1976):

$$\Phi_i(t_i) = \{Q(\cdot|t_i) \in \mathcal{P}_i(t_i) : Q \in \arg \max_{P \in \mathcal{P}_i} P(\tau_i^{-1}(t_i))\} \text{ for all } t_i \in T_i. \quad (3)$$

Thus, $\Phi_i(t_i)$ is the set of posteriors derived from priors that evaluates t_i as the most likely signal.

2.4 Definition of games and decision rules

To sum up the setup, an incomplete information game with multiple priors is a tuple $\mathcal{G} := \langle \Omega, \mathcal{I}, (\tau_i)_{i \in \mathcal{I}}, (\mathcal{P}_i)_{i \in \mathcal{I}}, (\Phi_i)_{i \in \mathcal{I}}, (A_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}} \rangle$. The incompleteness of information is expressed by the priors $(\mathcal{P}_i)_{i \in \mathcal{I}}$, the signals $(\tau_i)_{i \in \mathcal{I}}$, and the updating rules $(\Phi_i)_{i \in \mathcal{I}}$.

The equilibrium concepts we introduce in the next section adopt the following decision rules in \mathcal{G} . After t_i is observed, player i uses posteriors in $\Phi_i(t_i)$ to evaluate his actions. The interim payoff to a randomized action $\mu_i \in \Delta(A_i)$ given $\sigma_{-i} \in S_{-i} (= \prod_{j \neq i} S_j)$ and $Q \in \Phi_i(t_i)$ is

$$U_i(\mu_i, \sigma_{-i}|Q) = \sum_{\omega \in \Omega} \sum_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \mu_i(a_i) u_i(a_i, a_{-i}, \omega) \sigma_{-i}(a_{-i}|\tau_{-i}(\omega)) Q(\omega).$$

We write $U_i(a_i, \sigma_{-i}|Q)$ instead of $U_i(\mu_i, \sigma_{-i}|Q)$ if $\mu_i(a_i) = 1$. Since the set of actions and the set of states are finite, $U_i(\mu_i, \sigma_{-i}|Q)$ is continuous in (μ_i, σ_{-i}, Q) . We assume that each player uses an extremely pessimistic decision rule. That is, given the updated priors, we require that each player evaluates his actions using the worst possible scenario. Formally,

³See Gilboa and Schmeidler (1993) for axiomatization of updating rules.

after $t_i \in T_i$ is observed, the interim payoff to a randomized action $\mu_i \in \Delta(A_i)$ given $\sigma_{-i} \in S_{-i}$ is

$$V_i(\mu_i, \sigma_{-i}|t_i) = \min_{Q \in \Phi_i(t_i)} U_i(\mu_i, \sigma_{-i}|Q). \quad (4)$$

Notice that the interim payoff function is well behaved, continuous and concave in μ_i because U_i is continuous and the set $\Phi_i(t_i)$ is compact by assumption. The concavity follows since it is the minimum of linear functions of μ_i . But $V_i(\mu_i, \sigma_{-i}|t_i)$ is not necessarily linear in μ_i . So a player may strictly prefer to randomize actions, which will lead us to consider two different equilibrium concepts. Such an extreme decision rule is well studied in the decision theory literature (Gilboa and Schmeidler, 1989) and hence it constitutes one of natural specifications of games with uncertainty and ambiguity, but certainly not the only one.⁴

Notice that if each player can commit to his strategy ex ante, there is a natural ex ante equilibrium where player i 's payoff to strategy profile σ is simply

$$\min_{P \in \mathcal{P}_i} \sum_{\omega \in \Omega} \sum_{a \in A} u_i(a, \omega) \sigma(a|\tau(\omega)) P(\omega).$$

But as is discussed in the introduction, we are primarily interested in interim concepts, and the ex ante equilibrium will be referred in so far as to examine the implication of dynamic inconsistent behavior.⁵

3 Equilibrium concepts and existence

3.1 Equilibrium with multiple priors I: mixed strategy

We start with an equilibrium concept for $\mathcal{G} = \langle \Omega, \mathcal{I}, (\tau_i)_{i \in \mathcal{I}}, (\mathcal{P}_i)_{i \in \mathcal{I}}, (\Phi_i)_{i \in \mathcal{I}}, (A_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}} \rangle$ adopting the standard interpretation of mixed strategy in incomplete information games.

Definition 1 A strategy profile $\sigma^* \in S$ is a *mixed equilibrium* of \mathcal{G} if, for each $i \in \mathcal{I}$ and $t_i \in T_i$,

$$V_i(\sigma_i^*(t_i), \sigma_{-i}^*|t_i) \geq V_i(\mu_i, \sigma_{-i}^*|t_i) \quad (5)$$

for all $\mu_i \in \Delta(A_i)$.

That is, in a mixed equilibrium of \mathcal{G} , each player is maximizing his interim payoffs by choosing a lottery conditional on his signal. It is clear that if the set of priors \mathcal{P}_i is a singleton for all $i \in \mathcal{I}$, the mixed equilibrium of \mathcal{G} is equivalent to the standard Bayesian Nash equilibrium.

⁴It is known that, for some class of sets of priors, the decision rule of Gilboa and Schmeidler (1989) coincides with the decision rule based upon the Choquet integral (Schmeidler, 1989) with respect to convex capacities. In that case, our model of an incomplete information game with multiple priors can be defined as an incomplete information game with convex capacities.

⁵Obviously, if one never refers to the ex ante stage, one could make the collection of sets $\{\Phi_i(t_i) : t_i \in T_i\}$ as the sole primitive and the sets of priors are redundant. On the other hand, the model then could not discuss dynamic consistency and updating within the model, among others. We adopted this formulation since we are to model a game, not just particular equilibrium concepts.

The condition (5) trivially implies that $V_i(\sigma_i^*(t_i), \sigma_{-i}^*|t_i) \geq V_i(a_i, \sigma_{-i}^*|t_i)$ for any action $a_i \in A_i$, i.e., no pure action yields a higher payoff to player i . Since the interim payoff function $V_i(\mu_i, \sigma_{-i}|t_i)$ given by (4) is concave in μ_i , pure actions will often be strictly dominated by optimally mixed actions.⁶

A mixed equilibrium is a Nash equilibrium of the game where each t_i is treated as an independent player, and so under our finiteness assumption,⁷ it exists.

Proposition 2 *A mixed equilibrium of \mathcal{G} exists.*

Proof. We apply the standard existence theorem for a Nash equilibrium. Since Ω and A_i are finite, S_i is compact and convex for all $i \in \mathcal{I}$. For $\sigma_{-i} \in S_{-i}$, let $B_i(\sigma_{-i}) \subseteq S_i$ be the set of best responses of player i :

$$B_i(\sigma_{-i}) = \bigcap_{t_i \in T_i} \left\{ \sigma_i \in S_i : \sigma_i(t_i) \in \arg \max_{\mu_i \in \Delta(A_i)} V_i(\mu_i, \sigma_{-i}|t_i) \right\}.$$

We are done if $\sigma_{-i} \mapsto B_i(\sigma_{-i})$ is a non-empty, compact and convex valued, and upper hemicontinuous correspondence by applying Kakutani fixed point theorem to $\sigma \mapsto \prod_{i \in \mathcal{I}} B_i(\sigma_{-i})$.

Note that $U_i(\mu_i, \sigma_{-i}|Q)$ is continuous in (μ_i, σ_{-i}, Q) and that $\Phi_i(t_i)$ is compact. Since $V_i(\mu_i, \sigma_{-i}|t_i)$ is the minimum of $U_i(\mu_i, \sigma_{-i}|Q)$ over $Q \in \Phi_i(t_i)$, $V_i(\mu_i, \sigma_{-i}|t_i)$ is continuous in (μ_i, σ_{-i}) and concave in μ_i . Thus, for each t_i , the correspondence which maps σ_{-i} to the set $\{\sigma_i \in S_i : \sigma_i(t_i) \in \arg \max_{\mu_i} V_i(\mu_i, \sigma_{-i}|t_i)\}$ is non-empty, compact and convex valued, and upper hemicontinuous. Thus, $B_i(\sigma_{-i})$ is compact and convex valued, and upper hemicontinuous as the intersection of such correspondences. Finally, $B_i(\sigma_{-i})$ is non-empty, since for each t_i , $\{\sigma_i \in S_i : \sigma_i(t_i) \in \arg \max_{\mu_i} V_i(\mu_i, \sigma_{-i}|t_i)\}$ is non-empty and this set puts no restriction on the component corresponding to $t'_i \neq t_i$. This completes the proof. ■

Since a mixed equilibrium is a Nash equilibrium of a strategic form game, it inherits the standard properties of the Nash equilibrium. For instance, in a mixed equilibrium, no player ever uses a dominated action.

3.2 Equilibrium with multiple priors II: pure strategy

Allowing mixed strategy as in the previous subsection is a technically natural extension of the standard Bayesian Nash equilibrium. However, since preferences exhibit non-linearity in probability, the concept of mixed strategy is less innocuous than in the standard case. For instance, since the preference over mixed strategies is concave, the player may wish to randomize two actions if they are equally favorable. Then it is not clear what prevents the player from randomizing, beyond the strategy space specified before. Such possibilities are simply assumed away in the previous setup.

Thus we introduce an alternative notion, which is a natural analogue to the equilibrium in beliefs by Crawford (1990), so we shall adopt the same terminology.

⁶Such an example will be discussed in Subsection 4.3.

⁷Our existence results rely on the finiteness of the state space. See Section 5.4

Definition 3 A strategy profile $\sigma^* \in S$ is an *equilibrium in beliefs* of \mathcal{G} if, for each $i \in \mathcal{I}$ and $t_i \in T_i$, $\sigma_i(a_i|t_i) > 0$ implies

$$V_i(a_i, \sigma_{-i}^*|t_i) \geq V_i(a'_i, \sigma_{-i}^*|t_i) \quad (6)$$

for all $a'_i \in A_i$.

The equilibrium in beliefs of \mathcal{G} can be understood just as an equilibrium in beliefs for the complete information games in strategic form. In particular, in an equilibrium in beliefs σ^* , each player i is taking a pure action, but is believed to be randomizing over pure actions that are indifferent, as prescribed in σ_i^* , by the other players. Such beliefs are consistent with player i 's interim payoff maximization behavior, although it is not necessarily fully self-fulfilling in the sense that players' beliefs coincide with players' actual (methods of) choices of actions.

It is clear that if the set of priors \mathcal{P}_i is a singleton for all $i \in \mathcal{I}$, an equilibrium in beliefs of \mathcal{G} is equivalent to the standard Bayesian Nash equilibrium. It follows directly from the definitions that if a mixed equilibrium σ^* has the property that every player i at any t_i chooses a pure action, then it is an equilibrium in beliefs. When the updating rule is singleton-valued, then both equilibrium concepts coincide, since after updating the players' preferences are linear in probability assigned to actions. But in general, an equilibrium in beliefs of \mathcal{G} is not necessarily a mixed equilibrium of \mathcal{G} , nor vice versa.

The existence can be established by modifying Crawford's existence argument, as follows.⁸

Proposition 4 *An equilibrium in beliefs of \mathcal{G} exists.*

Proof. For each $i \in \mathcal{I}$, and for any $\sigma_{-i} \in S_{-i}$, let $B_i(\sigma_{-i})$ be defined by the rule:

$$B_i(\sigma_{-i}) = \bigcap_{t_i \in T_i} \left\{ \sigma_i \in S_i : \sigma_i(a_i|t_i) = 0 \text{ if } a_i \notin \arg \max_{a'_i \in A_i} V_i(a'_i, \sigma_{-i}|t_i) \right\}.$$

By construction, $B_i(\sigma_{-i})$ is non-empty. It is convex valued and upper hemicontinuous as the intersection of convex valued and upper hemicontinuous correspondences (note that $V_i(a'_i, \sigma_{-i}|t_i)$ is continuous in σ_{-i}). So the correspondence $\sigma \mapsto \prod_{i \in \mathcal{I}} B_i(\sigma_{-i})$ has a fixed point σ^* . Then by the construction of B_i , σ^* constitutes an equilibrium in beliefs of \mathcal{G} . ■

Since each player chooses a best action given information and the others' strategies, no player ever chooses an action which is dominated by another (pure) action. An equilibrium action may be dominated by a "mixed action" but like Crawford's idea in complete information games, the basic hypothesis here is that the players never randomize, and it is their beliefs which are in equilibrium.

⁸Again, the finiteness assumption is important. See Section 5.4 for discussion.

4 Examples: uncertainty under strategic interaction

We shall present examples which clarify the role of uncertainty or ambiguity in our model.

4.1 Difference of uncertainty induced by private information

In our setup, incomplete information can be expressed by differences in private information, as Harsanyi's Bayesian game. We shall give an example in which differences of uncertainty can be expressed by differences in private information.

Let there be two players, and consider a state space,

$$\Omega = \{1, 2, 3a, 3b, 4a, 4b\}$$

where the players have assigned probability $\varepsilon/2$ to the events $\{1\}$ and $\{2\}$ and probability $(1 - \varepsilon)/2$ to the events $\{3a, 3b\}$ and $\{4a, 4b\}$, respectively, where $0 < \varepsilon \leq 1$ is a given parameter. The difference between state $3a$ and state $3b$ and that between $4a$ and $4b$ are ambiguous in the sense that the players do not know how the probabilities assigned to $\{3a, 3b\}$ and $\{4a, 4b\}$ should be allocated to these states. Thus the players have a common set of priors, which is:

$$\mathcal{P}_1 = \mathcal{P}_2 = \left\{ P \in \Delta(\Omega) : P(\{1\}) = P(\{2\}) = \frac{\varepsilon}{2}, P(\{3a, 3b\}) = P(\{4a, 4b\}) = \frac{(1 - \varepsilon)}{2} \right\}.$$

Let $E = \{1, 2\}$. The following table summarizes the actions and payoffs, where Player 1 chooses a row and Player 2 chooses a column, and the numbers on the left are Player 1's payoffs and on the right are Player 2's payoffs.

| | $\omega \in E$ | | $\omega \notin E$ | |
|----------|----------------|---------|-------------------|---------|
| | α | β | α | β |
| α | 1, -2 | 0, 0 | 1, 1 | 0, 0 |
| β | 0, -2 | 1, 0 | 0, 1 | 1, 0 |

Note that Player 1's best response is to choose the action Player 2 chooses. Given ω , Player 2's payoffs are independent of Player 1's choice of action, and Player 2 wants to choose β if $\omega \in E$ and to choose α if $\omega \notin E$.

As a bench mark, consider the case where there is no private information. Then both players agree that event E occurs with probability ε , and this is common knowledge. So in this game the multiplicity of priors is inessential if both players remain uninformed about the state. If ε is small enough, then playing α maximizes Player 2's payoff regardless of Player 1's behavior, and thus a unique equilibrium is that both players choose α .

Now let us consider private information. The ranges of signals (τ_1, τ_2) are

$$T_1 = \{\{1, 3a, 3b\}, \{2, 4a, 4b\}\},$$

$$T_2 = \{\{1, 3a, 4a\}, \{2, 3b, 4b\}\},$$

where $\tau_i(\omega) \in T_i$ is the set containing $\omega \in \Omega$. The reader may find it easy to think of this as if the private information is given by partitions suggested by T_1 and T_2 .

We assume that both players use the FB updating rule (2). So the set of updated priors is that of all probability distributions which are consistent with observation. We have:

$$\begin{aligned}\Phi_1(\{1, 3a, 3b\}) &= \{P \in \Delta(\Omega) : P(\{1\}) = \varepsilon, P(\{3a, 3b\}) = 1 - \varepsilon\}, \\ \Phi_1(\{2, 4a, 4b\}) &= \{P \in \Delta(\Omega) : P(\{2\}) = \varepsilon, P(\{4a, 4b\}) = 1 - \varepsilon\}, \\ \Phi_2(\{1, 3a, 4a\}) &= \left\{ P \in \Delta(\Omega) : P(\{1\}) = \frac{\varepsilon}{\varepsilon + 2(x+y)}, P(\{3a\}) = \frac{2x}{\varepsilon + 2(x+y)}, \right. \\ &\quad \left. P(\{4a\}) = \frac{2y}{\varepsilon + 2(x+y)} \text{ where } x \in \left[0, \frac{1-\varepsilon}{2}\right], y \in \left[0, \frac{1-\varepsilon}{2}\right] \right\}, \\ \Phi_2(\{2, 3b, 4b\}) &= \left\{ P \in \Delta(\Omega) : P(\{2\}) = \frac{\varepsilon}{\varepsilon + 2(x+y)}, P(\{3b\}) = \frac{2x}{\varepsilon + 2(x+y)}, \right. \\ &\quad \left. P(\{4b\}) = \frac{2y}{\varepsilon + 2(x+y)} \text{ where } x \in \left[0, \frac{1-\varepsilon}{2}\right], y \in \left[0, \frac{1-\varepsilon}{2}\right] \right\}.\end{aligned}$$

We shall find a unique equilibrium of this game. Note that the updated probabilities of E are:

$$\{P(E) | P \in \Phi_1(t_1)\} = \{\varepsilon\}, \quad \{P(E) | P \in \Phi_2(t_2)\} = [\varepsilon/(2-\varepsilon), 1]$$

for all $t_1 \in T_1$ and $t_2 \in T_2$.⁹ Thus, within each player, the evaluation of E is the same for every state $\omega \in \Omega$. But in spite that the players start with a common set of multiple priors, they have different uncertainty concerning E when $\varepsilon > 0$; Player 1 assigns the unique probability ε and Player 2 assigns multiple probabilities ranging from $\varepsilon/(2-\varepsilon)$ to 1. The difference of uncertainty is attributed to the difference of private information T_1 and T_2 .

To find optimal actions, let p be the probability Player 2 chooses action α . For any $t_2 \in T_2$, the interim payoff of Player 2 is

$$\begin{aligned}\min_{P \in \Phi_2(t_2)} & ((-2p + 0 \cdot (1-p)) \cdot P(E) + (1p + 0 \cdot (1-p)) \cdot (1 - P(E))) \\ &= \min_{P(E) \in [\varepsilon/(2-\varepsilon), 1]} (-3pP(E) + p) = -3p + p = -2p\end{aligned}$$

if $p > 0$, and it is 0 if $p = 0$, which implies that β is a strictly dominant action for all $t_2 \in T_2$. Knowing this, Player 1, who wants to match his action, must choose β for all $t_1 \in T_1$. To summarize, the game has a unique equilibrium, in both definitions we proposed, in which both players always choose β . Note that this is true for any small $\varepsilon > 0$.

Now look at the case $\varepsilon = 0$ where there is no difference of uncertainty concerning the event E because E is null:

$$\{P(E) | P \in \Phi_1(t_1)\} = \{P(E) | P \in \Phi_2(t_2)\} = \{0\}.$$

Thus payoffs are given by the table corresponding to $\omega \notin E$ with probability one. So we have a complete information game with a randomization device τ , but since α is a strictly

⁹Though players have different uncertainty, there is a common probability in the sense that $\{P(E) | P \in \Phi_1(t_1)\} \cap \{P(E) | P \in \Phi_2(t_2)\} \neq \emptyset$. We will discuss this issue in the next section.

dominant action for Player 2, we conclude that there is a unique equilibrium where both players choose α .

In conclusion, the equilibrium set changes discontinuously with respect to ε at $\varepsilon = 0$. Note that even at $\varepsilon = 0$, the players have multiple priors. But Player 2's set of updated probabilities of the relevant event E gets degenerate at $\varepsilon = 0$, and this fact generates the discontinuity. Intuitively, Player 2, being very pessimistic, hesitates to choose action α which is very bad when E occurs, as long as there is some chance that E is true. When $\varepsilon = 0$, his worry disappears and he is willing to choose α .

Notice the dynamic inconsistency of Player 2's behavior: suppose Player 2 could commit in advance to the action he will be playing after arrival of private information. That is, we look at an ex ante equilibrium of the game. Then for any prior in \mathcal{P}_2 , the value of committing to always playing α is $-2 \times \varepsilon + 1 \times (1 - \varepsilon)$, which is positive if ε is small, better than the value of committing to β whose value is zero. Since Player 1 is just reacting to Player 2's action, the difference of the equilibrium behavior from the bench mark case can also be attributed (at least partially) to the dynamic inconsistent behavior of Player 2.

But for a game theoretic implication, a more important point is the way Player 1's behavior is affected. Notice that Player 1 unambiguously assigns a single probability to the payoff relevant event E , irrespective of the value of ε and his private information. So as far as his payoffs are concerned, he has no uncertainty at all. But Player 1 knows that Player 2 tends to interpret his private information very pessimistically, and Player 1 must take this into account in equilibrium. In fact, this example can be modified that Player 1 has a single prior; just take any prior in \mathcal{P}_1 . Hence this example can also be seen as an instance of a standard Bayesian player's action is affected by the other non-Bayesian player.

4.2 An equilibrium with no BNE justification and the role of updating rule

We shall give a simple example where there is an equilibrium with multiple priors under the FB updating rule (2) which cannot be justified as a Bayesian Nash equilibrium for the given state space and information structure, and moreover it is not an equilibrium if the ML updating rule (3) is adopted. So this is also an example to understand the role of updating as well as multiplicity of priors.

We use the same notation as in the previous example. Let $\Omega = \{1, 2a, 2b, 3\}$ with

$$\mathcal{P}_1 = \mathcal{P}_2 = \{P \in \Delta(\Omega) : P(1) = 0.25, P(\{2a, 2b\}) = 0.5, P(3) = 0.25\}.$$

Let the payoffs be given by the following table.

| | | $\omega = 1$ | | $\omega \in \{2a, 2b\}$ | | $\omega = 3$ | | | |
|----------|--|--------------|---------|-------------------------|---------|--------------|---------|------|------|
| | | α | β | α | β | α | β | | |
| α | | 3, 3 | 0, 1 | | 1, 3 | 1, 1 | | 3, 3 | 0, 1 |
| β | | 0, 0 | 1, 1 | | 0, 0 | 0, 1 | | 0, 0 | 1, 1 |

The signals are given by:

$$T_1 = \{\{1\}, \{2a\}, \{2b\}, \{3\}\}, \quad T_2 = \{\{1, 2a\}, \{2b, 3\}\}.$$

Player 1 always knows the payoffs, and Player 2's payoffs do not depend on ω . So there is no "uncertainty" in the payoff structure.

With the FB updating rule, the following strategy profile constitutes an equilibrium with multiple priors for both definitions.

$$\begin{aligned} \sigma_1(\{1\}) &= \beta, & \sigma_1(\{2a\}) &= \alpha, & \sigma_1(\{2b\}) &= \alpha, & \sigma_1(\{3\}) &= \beta, \\ \sigma_2(\{1, 2a\}) &= \beta, & & & \sigma_2(\{2b, 3\}) &= \beta. \end{aligned}$$

Let us confirm this. Since α is a dominant action for Player 1 when $\omega \in \{2a, 2b\}$, playing α in $2a$ and $2b$ is optimal. If Player 2 plays β , β is a best response for Player 1 when $\omega \in \{1, 3\}$. So Player 1's behavior is optimal. Player 2's behavior can be shown to be optimal by a similar calculation as in the previous example. Intuitively, when Player 1 is to play β in state 1, after observing $\{1, 2a\}$, if Player 2 is to play α with some probability, he will assign probability one to state 1 which is the worst scenario, and then playing β for sure is a best response. A symmetric argument applies for $\{2b, 3\}$.

But the strategy profile given above cannot be a Bayesian Nash equilibrium for any single prior $P \in \mathcal{P}_i$. The reason is as follows: since $P(2a) + P(2b) = 0.5$ must hold, one of $2a$ and $2b$ must have ex ante probability no less than 0.25. Assume that it is $2a$ without loss of generality since the structure of the game is symmetric. Then after $\{1, 2a\}$ is observed, Player 2 must assign at least probability 0.5 to state $2a$. Since Player 1 plays α in state $2a$, this implies that Player 2 knows that α is played at least probability 0.5, and then he must choose α since it is the risk dominant action for Player 2.

The strategy profile above is not an equilibrium with the ML updating rule. For Player 2, after observing $\{1, 2a\}$, the prior which makes this most likely is the one assigning ex ante probability of 0.5 to $2a$. Thus he believes action β and α occurs with ratio 0.25 : 0.5, and so Player 2 must play α . Since it is never a best response of Player 1 to play β with any probability in states $2a$ and $2b$, we see that Player 2 must always play α in any equilibrium. Then Player 1 must always play α since it is a best response to α for any ω . In conclusion, with the ML updating rule, in a unique equilibrium strategy profile, both players always play α .

4.3 Difference of two equilibrium concepts

Let us give an example to clarify the differences of the two notions of equilibria we proposed.¹⁰ Let $\mathcal{I} = \{1, 2\}$ and $\Omega = \{1, 2\}$. The set of priors $\mathcal{P}_1 = \mathcal{P}_2$ is just the set of all probability distributions on Ω . Player 2 observes every $\omega \in \Omega$, but Player 1 has no private information. Player 1 has three actions, α, β, γ and Player 2 has two actions 1 and 2.

Independent of Player 1's action, Player 2's payoff is 10 if he chooses action ω when $\omega \in \Omega$ occurs, otherwise 0. Then for both equilibrium concepts, Player 2 must play action

¹⁰This example and discussion are inspired by Lo (1996).

ω when $\omega \in \Omega$ is observed. This makes the game effectively a single person decision problem of Player 1 against the nature, where the decision maker has multiple priors \mathcal{P}_1 . The payoffs of Player 1 are as follows.

| | | |
|----------|----|----|
| | 1 | 2 |
| α | 10 | 0 |
| β | 0 | 10 |
| γ | 1 | 1 |

Then the payoffs to actions α and β are both zero owing to Player 1's pessimistic expectations, whereas action γ is worth 1. So a unique equilibrium in beliefs is that Player 1 chooses γ and Player 2 behaves as described above. But if mixed actions are considered, randomizing equally between α and β yields 5 irrespective of priors, and this is payoff maximizing. So in a unique mixed equilibrium, Player 1 chooses this randomization strategy.

5 Discussions and generalizations

5.1 Foundations

Let us comment on a couple of issues on the foundation of our model: preferences and types.

For preferences, we simply postulated that there exist sets of priors and that the players' preferences over strategies are induced by the most pessimistic posteriors. A possible objection is lack of axiomatic foundations, and our defense is as follows. As for the use of multiple priors, in a single person setting, there is a well known axiomatization of such preference relations by Gilboa and Schmeidler (1989). Even for complete information games, a common justification for expected utility preferences is based on individual decision making. So if this line of justification is accepted, we contend that the Gilboa-Schmeidler axiomatization justifies our multiple priors approach as well.

To discuss the role of types in our model, let us first review the standard approach: each point of a type space are associated a state of nature as well as a single posterior for each player on the type space itself. The type of a player encodes not only his beliefs on the space, but also his beliefs about others' beliefs, his beliefs about others' beliefs about others' beliefs, and so on. Whether or not this entire sequence of beliefs can be captured in a single type space is a fundamental question to the construction and analysis of games with incomplete information. Mertens and Zamir (1985) did the first mathematical analysis on this, and gave an affirmative answer.

In our framework, we postulate a multiple posterior version of type spaces as above: with each point of a type space are associated a state of nature as well as multiple posteriors for each player on the type space itself. Then a similar question will naturally arise. Can one construct a type space with multiple priors from a hierarchy of sets of beliefs?

A seminal work in this direction is Epstein and Wang (1996), which criticize the Bayesian approach of comprehensive belief types and provide a preference based construc-

tion of general type spaces which serves as a foundation for games with incomplete information. Our model conforms to their general definition of incomplete information games. But the preference based approach, by construction, does not spell out the structure of sets of beliefs independently.

More closely related to our model is Ahn (2003)’s construction of type spaces. He considered players with multiple priors over the state space, which is similar to our model, and studied generalization of the classic coherency condition which produces a Mertens-Zamir style type space to encode the hierarchy of beliefs. He discussed a class of incomplete information games with multiple priors to demonstrate that his result serves as its foundation. The class of games in Ahn (2003) consist of players with multiple priors and general preferences over players’ beliefs and consequences. It can be seen that if the preferences are defined in terms of the decision rule of Gilboa and Schmeidler (1989), the class of games are reduced to our model. In this sense, Ahn (2003)’s result directly serves as a foundation of our model.¹¹

5.2 Common priors and “agreeing to disagree”

In the single prior model, it has been pointed out that the common prior assumption has excessively strong economic implications; e.g., various no trade results originated in Milgrom and Stokey (1982). This excess power of the model gives rise to some skepticism, and indeed it has motivated many works on relaxing the common prior assumption (more importantly common support assumption) in the literature. So from this point of view, it is desirable to consider a model allowing for different sets of priors, which we chose to do.

On the other hand, it is not necessarily clear if our model exhibits similar drawbacks if we assume a common set of priors. As a matter of fact, the implications we draw from the examples in Section 4 do not rely on the heterogeneity of sets of priors at all because all the examples assume a common set of priors.

Since there are two channels, risks and uncertainty, through which private information makes differences, our model is much richer than the single prior model, even if a common set of priors is assumed in addition. For instance, we saw in Subsection 4.2 a game with common multiple priors where there is an equilibrium which cannot be explained by a single common prior. Therefore, even if a common set of priors is required in addition, it does not seem to be too restrictive as far as our main messages are concerned.

It is an interesting future research topic to study whether a common set of priors is a theoretically natural requirement. To make a contribution on this direction, we investigate implications of a common set of priors, by considering the role of “common knowledge” via the problem of “agreeing to disagree” à la Aumann (1976) in our setup.

Since our premise is that the structure of the game is completely determined by state $\omega \in \Omega$, and the private information of player i is given by τ_i , it is natural to adopt

¹¹Ahn (2003) also defined an equilibrium concept, which can be translated in our framework. Roughly speaking, a strategy profile is an equilibrium in Ahn’s sense in our model if it is both a mixed equilibrium and a equilibrium in beliefs in our model. Since there are games where these two concepts do not imply each other as we saw in Subsection 4.3, Ahn’s equilibrium may not exist in general.

Aumann's formulation of common knowledge in our setup. That is, an event $E \subseteq \Omega$ is common knowledge (at $\omega \in E$) if E contains a partition element of the finest common coarsening of partitions $\{\tau_i^{-1}(t_i) : t_i \in T_i\}$, $i = 1, \dots, I$. Thus the multiplicity of priors does not play any role in determining whether or not an event is common knowledge among the players.

Let an event $E \subseteq \Omega$ be given. For each player $i \in \mathcal{I}$ and a state $\omega \in \Omega$, define $\rho_i(E|\omega)$ by the rule:

$$\rho_i(E|\omega) = \{p \in [0, 1] : p = P(E), P \in \Phi_i(\tau_i(\omega))\},$$

which is the collection of player i 's ex post evaluation of E at $\omega \in \Omega$.

A natural question in view of Aumann's theorem is if players with common multiple priors can agree to disagree, and if they do, to what extent they agree. We attempt to answer this question by the following proposition, which contains Aumann's agreement theorem as a special case.

Proposition 5 *Let $E \subseteq \Omega$ be an event. Suppose that*

- $\mathcal{P}_i = \mathcal{P} \subseteq \Delta(\Omega)$ for all $i \in \mathcal{I}$,
- every player adopts the FB updating rule,
- for all $i \in \mathcal{I}$, $\rho_i(E|\omega)$ is common knowledge at $\omega \in \Omega$: that is, the event $\{\omega' \in \Omega : \rho_i(E|\omega) = \rho_i(E|\omega')\}$ is common knowledge,
- for all $i \in \mathcal{I}$, $\rho_i(E|\omega)$ is a closed interval.

Then we have:

$$\bigcap_{i \in \mathcal{I}} \rho_i(E|\omega) \neq \emptyset.$$

Proof. The common knowledge assumption implies that there exists a common knowledge event F with $\omega \in F$ such that $\rho_i(E|\omega') = \rho_i(E|\omega)$ for all $\omega' \in F$ for each $i \in \mathcal{I}$. For any $\omega' \in F$, set $t_i = \tau_i(\omega')$, and then it must be true that $\tau_i^{-1}(t_i) \subseteq F$ since F is common knowledge. Thus, $P(E|t_i) = P(\tau_i^{-1}(t_i) \cap E) / P(\tau_i^{-1}(t_i)) = P(\tau_i^{-1}(t_i) \cap F \cap E) / P(\tau_i^{-1}(t_i)) = P(E \cap F|t_i)$. That is, the conditional probabilities of E and $E \cap F$ are the same at any $\omega' \in F$ for any $P \in \mathcal{P}$.

Let $p_*(i), p^*(i) \in [0, 1]$ be the upper bound and the lower bound of $\rho_i(E|\omega')$:

$$p_*(i) = \min_{p \in \rho_i(E|\omega')} p = \min_{P \in \mathcal{P}} P(E|\tau_i(\omega')) = \min_{P \in \mathcal{P}} P(E \cap F|\tau_i(\omega')),$$

$$p^*(i) = \max_{p \in \rho_i(E|\omega')} p = \max_{P \in \mathcal{P}} P(E|\tau_i(\omega')) = \max_{P \in \mathcal{P}} P(E \cap F|\tau_i(\omega')),$$

for all $\omega' \in F$.

Let $P_*, P^* \in \mathcal{P}$ be such that:

$$P_*(E|F) = P_*(E \cap F)/P_*(F) = \min_{P \in \mathcal{P}} P(E \cap F)/P(F),$$

$$P^*(E|F) = P^*(E \cap F)/P^*(F) = \max_{P \in \mathcal{P}} P(E \cap F)/P(F).$$

Then, for any $t_i \in \tau_i(F)$, we have:

$$\begin{aligned} p_*(i) &\leq P_*(E \cap F | t_i) = \frac{P_*(E \cap F \cap \tau_i^{-1}(t_i))}{P_*(\tau_i^{-1}(t_i))}, \\ p^*(i) &\geq P^*(E \cap F | t_i) = \frac{P^*(E \cap F \cap \tau_i^{-1}(t_i))}{P^*(\tau_i^{-1}(t_i))}, \end{aligned}$$

or equivalently,

$$\begin{aligned} p_*(i)P_*(\tau_i^{-1}(t_i)) &\leq P_*(E \cap F \cap \tau_i^{-1}(t_i)), \\ p^*(i)P^*(\tau_i^{-1}(t_i)) &\geq P^*(E \cap F \cap \tau_i^{-1}(t_i)). \end{aligned}$$

Since both P_* and P^* are priors and so they are additive, summing the above over $t_i \in \tau_i(F)$, we have

$$\begin{aligned} p_*(i)P_*(F) &\leq P_*(E \cap F), \\ p^*(i)P^*(F) &\geq P^*(E \cap F). \end{aligned}$$

This means that $[P_*(E|F), P^*(E|F)] \subseteq \rho_i(E|\omega)$ for every i , so we have established the result. ■

This result is an extension of Aumann's theorem since if \mathcal{P}_i is singleton, so is each $\rho_i(E|\omega)$ and hence $\bigcap_{i \in \mathcal{I}} \rho_i(E|\omega) \neq \emptyset$ implies the agreement $\rho_i(E|\omega) = \rho_j(E|\omega)$ for any $i, j \in \mathcal{I}$.

We can interpret $\bigcap_{i \in \mathcal{I}} \rho_i(E|\omega) = \emptyset$ as a situation where the players completely disagree about posterior beliefs on E . So the proposition implies that complete disagreement cannot be common knowledge, and so it has the flavor of Aumann's theorem: if posteriors are common knowledge, they must share a posterior belief as one of the posterior beliefs in $\rho_i(E|\omega)$.

In general, the sets $\rho_i(E|\omega)$, $i = 1, \dots, I$, differ from each other, and thus there are posterior beliefs in their posterior belief sets which do not belong to each others' posterior belief set. Consider the example discussed in Subsection 4.1. We have already obtained

$$\rho_1(E|\omega) = \{\varepsilon\}, \quad \rho_2(E|\omega) = [\varepsilon/(2 - \varepsilon), 1]$$

for all $\omega \in \Omega$. Thus, $\rho_1(E|\omega)$ and $\rho_2(E|\omega)$ are common knowledge at every $\omega \in \Omega$. Clearly, $\rho_1(E|\omega) \neq \rho_2(E|\omega)$ and $\rho_1(E|\omega) \cap \rho_2(E|\omega) = \rho_1(E|\omega)$.

The problem of "agreeing to disagree" is closely related to various no trade results, and this issue for multiple prior models is discussed in Kajii and Ui (2004). They present a framework to understand the possibility of a purely speculative trade under asymmetric information, where the decision making rule of each trader conforms to the multiple priors model. In this framework, they derive a necessary and sufficient condition on the sets of posteriors, thus implicitly on the updating rules adopted by the players, for non-existence of trade such that it is always common knowledge that every player expects a positive gain. As a corollary of the main result, they obtain generalization of Proposition 5, which states that, not only when $\mathcal{P}_i = \mathcal{P}_j$ for all $i, j \in \mathcal{I}$, but also when $\bigcap_{i \in \mathcal{I}} \mathcal{P}_i \neq \emptyset$, we obtain $\bigcap_{i \in \mathcal{I}} \rho_i(E|\omega) \neq \emptyset$. For more details, see Kajii and Ui (2004).

5.3 Complete information games

Consider the special cases where payoff functions are independent of state ω . Thus signals and priors can be seen as external randomization devices. Then our model can be understood as a complete information game with an external randomization device, but there is uncertainty about which randomization devices are to be used. Let us relate this to the studies on generalizations of Nash equilibrium for complete information games, by allowing uncertainty averse players. These studies include Dow and Werlang (1994), Eichberger and Kelsey (2000), Klibanoff (1996), Marinacci (2000), and Lo (1996, 2002). It can be shown that the mixed equilibrium in our definition is a correlated equilibrium under uncertainty in Lo (2002) in general, and when the signals are independent for any priors, it corresponds to the equilibrium (with agreement) of Lo (1996).

In the studies of the generalized Nash equilibrium, the issue is *endogenous* formation of beliefs, accommodating players' possibly non-additive beliefs on the other's actions. On the other hand, in our framework, player's beliefs about the opponents' action distributions are formed from two components: the first is about the opponents' strategic choices of actions *after* they have observed signals, and the second is inference about what private signals the opponents have observed, given the player's private information. The player's beliefs may be non-additive for the second components, but they are additive for the first. So the possibly non-additive part is *exogenous* in our framework, which makes our model very tractable.

Moreover, there is a conceptual difficulty of endogenously formed non-additive beliefs, in particular in dynamic settings with more than 2 players, that it is not clear how updating is done. In our setup the updating is controlled exogenously so there will be no conceptual issues of this kind.

Coming back to the example discussed in Subsection 4.3, Lo (1996) used essentially the same example to justify why he favors mixed equilibrium over equilibrium in beliefs. One of the reasons is that an action dominated by a mixed action may survive in equilibrium in beliefs with uncertainty aversion: a player may form an excessively pessimistic beliefs about the others' choices and so he may assign an excessively low value to something whose payoffs depend on the others' choices, even if he could perfectly hedge risks had he not been occupied with an extremely pessimistic view. Also, since the support of a non-additive measure is not necessarily well defined, the equilibrium in beliefs requires well defined support by definition.

However, we do not particularly favor the mixed equilibrium over the equilibrium in beliefs. First of all, both make good sense. Secondly, it is not necessarily a defect of the equilibrium concept if a particular player chooses a "dominated" action. In strategic environments, we are often interested in strategic implications of dominated action to *the other players*. Thirdly, in our model, the equilibrium in beliefs do not have this problem about the support. Roughly speaking, our equilibrium in beliefs can be regarded as a purified version of the equilibrium in beliefs à la Dow and Werlang (1994), and it is well behaved.

Since a standard Nash equilibrium is always an equilibrium with uncertainty aversion,

if endogenous formation of non-additive beliefs is considered, there is a large number of equilibria which include the ones supported by extremely pessimistic expectations of some players. Therefore, although the generalized Nash equilibrium concepts explain why players may be stuck in a situation by pessimism, they do not have any stronger predictive power than the standard Nash equilibrium. In order to be used for economic analysis, one need to think about “refinements” of those endogenous equilibria, as is done in Marinacci (2000) for instance, after all.

Our position is that the payoff matrix should be “complete” by itself and it is more natural to model any extra doubt and uncertainty about how players might have among themselves as an exogenous state space (and multiple priors) outside the payoff matrix. If one has to apply an extraneous refinement criteria on “equilibria” affected by uncertainty anyway, why not describe the uncertainty specifically in the model.

5.4 On finiteness assumptions and generalizations

Let us conclude by mentioning possible generalization of our models. We assumed the state space and action sets to be finite sets. But the definitions of equilibria we proposed are very general. For instance, an auction model of Lo (1998) conforms to our model except for finiteness.¹² It is worth pointing out that the so called global games which assumes an “improper prior” of uniform distribution over the entire real line¹³ can be readily interpreted as multiple priors with the ML updating rule in our model. Conceptually, there is not much reason for using finite models, and in fact one would need to use infinite state space to model complex structures of incomplete information.

We chose a finite model to make the existence results transparent, but more importantly it is one of our purposes that even a simple finite model is rich enough to explain interesting phenomena. Technically, we are confident that the existence results can be extended to allow an infinite state space and infinite actions once we assume strong enough continuity and compactness. The existence problem for a general infinite state space is delicate even for the standard Bayesian Nash equilibria, so our model naturally inherits the same difficulty, and in addition, the equicontinuity property of payoff functions indexed by the updated prior sets $(\Phi_i)_{i \in \mathcal{I}}$ is tricky in general spaces.

References

- [1] Ahn, D. (2003). “Hierarchies of Ambiguous Beliefs,” mimeo. Available from <http://www.stanford.edu/~ahnd/hierarchies.pdf>
- [2] Aumann, R. J. (1976). “Agreeing to Disagree,” *Annals of Statistics* **4**, 1236–1239.
- [3] Chen, Y., Katuscak, P., and Ozdenoren, E. (2003). “Sealed Bid Auctions with Ambiguity: Theory and Experiments,” mimeo. Available from http://www.si.umich.edu/~yanchen/papers/chen_katuscak_ozdenoren_20030711.pdf

¹²In Lo’s model, private signals are assumed to be *independent*, which makes his model very tractable.

¹³See a survey by Morris and Shin (2002).

- [4] Crawford, V. (1990). "Equilibrium without Independence," *Journal of Economic Theory* **50**, 127–154.
- [5] Dempster, A. P. (1967). "Upper and Lower Probabilities Induced by a Multi-valued Mapping," *Annals of Mathematical Statistics* **38**, 325–339.
- [6] Dow, J., and Werlang, S. R. C. (1994). "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory* **64**, 304–324.
- [7] Eichberger, J., and Kelsey, D. (2000). "Non-additive Beliefs and Strategic Equilibria," *Games and Economic Behavior* **30**, 183–215.
- [8] Epstein, L., and Wang, T. (1996). "Beliefs about Beliefs' without Probabilities," *Econometrica* **64**, 1343–1373.
- [9] Epstein, L., and Zhang, J. (2001). "Subjective Probabilities on Subjectively Unambiguous Events," *Econometrica* **69**, 265–306.
- [10] Fagin, R., and Halpern, J. (1990). "A New Approach to Updating Beliefs," *Proceedings of the 6th Conference on Uncertainty in AI*, 317–325.
- [11] Gilboa, I., and Schmeidler, D. (1989). "Maxmin Expected Utility with Non-unique Prior," *Journal Mathematical Economics* **18**, 141–153.
- [12] Gilboa, I., and Schmeidler, D. (1993). "Updating Ambiguous Beliefs," *Journal of Economic Theory* **59**, 33–49.
- [13] Harsanyi, J. (1967–68). "Games with Incomplete Information Played by Bayesian Players," Parts I, II, and III, *Management Science* **14**, 159–182, 320–334, 486–502.
- [14] Jaffray, J. -Y. (1992). "Bayesian Updating and Belief Function," *IEEE Transactions on Systems, Man and Cybernetics* **22**, 1144–1152.
- [15] Kajii, A., and Ui, T. (2004). "Trade with Heterogeneous Multiple Priors," KIER Working Paper 582, Kyoto University.
- [16] Klibanoff, P. (1996). "Uncertainty, Decision, and Normal Form Games," mimeo. Available from <http://www.kellogg.nwu.edu/faculty/klibanof/ftp/workpap/unc.pdf>
- [17] Lo, K. C. (1996). "Equilibrium in Beliefs under Uncertainty," *Journal of Economic Theory* **71**, 443–484.
- [18] Lo, K. C. (1998). "Sealed Bid Auctions with Uncertainty Averse Bidders," *Economic Theory* **12**, 1–20.
- [19] Lo, K. C. (2002). "Correlated Equilibrium under Uncertainty," *Mathematical Social Sciences* **44**, 183–209.

- [20] Marinacci, M. (2000). “Ambiguous Games,” *Games and Economic Behavior* **31**, 191–219.
- [21] Mertens, J. F., and Zamir, S. (1985). “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory* **14**, 1–29.
- [22] Milgrom, P., and Stokey, N. (1982). “Information, Trade, and Common Knowledge,” *Journal of Economic Theory* **26**, 17–27.
- [23] Ozdenoren, E. (2002). “Auctions and Bargaining with a Set of Priors,” mimeo, University of Michigan.
- [24] Salo, A., and Weber, M. (1995). “Ambiguity Aversion in First-Price Sealed-Bid Auctions,” *Journal of Risk and Uncertainty* **11**, 123–137.
- [25] Schmeidler, D. (1986). “Integral Representation without Additivity,” *Proceedings of the American Mathematical Society* **97**, 255–261.
- [26] Schmeidler, D. (1989). “Subjective probability and Expected Utility without Additivity,” *Econometrica* **57**, 571–587.
- [27] Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton Univ. Press.