

Equilibrium Departures from Common Knowledge in Games with Non-Additive Expected Utility

Sujoy Mukerji
Department of Economics
Oxford University

Hyun Song Shin
Department of Accounting and Finance
London School of Economics

Final Revised Version Dec 2001

Abstract

This paper concerns the interpretation of equilibrium in non-additive beliefs in two-player normal form games. We argue that such equilibria involve beliefs and actions which are consistent with a lack of common knowledge of the game. Our argument rests on representation results which show that different notions of equilibrium in games with non-additive beliefs may be reinterpreted as equilibrium in associated games of incomplete information with *additive* (Bayesian) beliefs where common knowledge of the (original) game does not apply. The representation results show one way of comparing and understanding the various notions of equilibrium, for games with non-additive beliefs, advanced in the literature.

Keywords: Knightian uncertainty, uncertainty aversion, ambiguity aversion, non-additive probabilities, capacities, belief functions, Choquet expectation, common belief, common knowledge, non-cooperative games.

JEL Classification: C72, C79, D80, D81

1 Introduction

The epistemic foundations of equilibrium in games has been one of the notable advances in game theory in recent years. This development has also coincided with the development of tools using non-additive probabilities and associated formal techniques to model behavior which has been labeled as ambiguity aversion, or uncertainty aversion. The environment which elicits such behavior has been dubbed “Knightian uncertainty”, following Knight’s (1921) distinction between risk and uncertainty. Building on the successful application of these methods to single-person decision theory, a burgeoning literature has attempted to apply them to games.¹

This paper concerns the interpretation of equilibrium in non-additive beliefs in *two*-player normal form games. We argue that such equilibria involve beliefs and actions which are consistent with a lack of common knowledge of the game. Our argument rests on representation results which show that different notions of equilibrium in games with non-additive beliefs may be reinterpreted as standard notions of equilibrium in associated games of incomplete information with *additive* (Bayesian) beliefs where common knowledge of the (original) game does not apply. We show that any pair of non-additive belief functions (and actions, to the extent these are explicit in the relevant notion of equilibrium) which constitute an equilibrium in the game with Knightian uncertainty may be replicated as beliefs and actions of a specific pair of *types*, one for each player, in an equilibrium of an orthodox Bayesian game, in which there is a common prior over the type space.

A principal inspiration for the construction of the incomplete information games invoked in the analysis in this paper was the missing/misspecified states interpretation of non-additive probabilities.² According to this interpretation, a DM’s non-additive belief over a set of “objectively” given states is a reduced form description of a standard (additive) probabilistic belief over a “fuller” state space. The fuller state space has, in addition to the “objective” states, a set of “subjective” states, which are further possibilities imagined by the DM. In this paper, we start with a *given* game (of complete information) and a *given* equilibrium in non-additive beliefs of the game. Then, *corresponding* to this given game and the given equilibrium, we construct an incomplete information game in which players may be one of several types.³ The posterior beliefs of a certain pair of types, in the constructed Bayesian game, is an additive representation of the non-additive beliefs in the equilibrium of the given game. The posterior beliefs are described on a state space that has more states than in the original game since that game, being a game of complete information, allows for just one type for each player. In the constructed Bayesian game, the additional states deemed possible by player 1 comprise of various perturbations of player 2’s payoffs and the possibility that player 2 may be able to choose after obtaining information about player 1’s choice. Player 2’s non-additive belief in the equilibrium of the given game is represented as an additive posterior on an analogously extended state space. The extended state space allows for perturbations of player 1’s payoffs and also of the information structure of the game. Each player knows his own type, i.e., knows his own payoffs, which are as in the original game. Hence,

¹Papers applying the non-additive expected utility framework to games include Dow and Werlang (1994), Groes et al. (1998), Haller (1995), Eichberger and Kelsey (2000), Epstein (1997), Mukerji (1995), Marinacci (2000) and Ghirardato and Le Breton (2000). Also, Klibanoff (1996), Lo (1996) and Epstein (1997) analyse games using the multiple priors model, which is closely related to the non-additive expected utility approach.

²The next section contains a more detailed discussion of the idea and the relevant citations.

³Recall, in an incomplete information game, a state is a specification of a pair of types, one for each player.

the supports of the players' posteriors (on the space of types) are common only in part; the types in the common part have the payoffs exactly as in the given game. The common p -belief of the actual game being played is generated by that portion of the support of posterior beliefs which is shared by the players. Thus, the given game is not common-1-belief in the incomplete information game. Indeed, we find that the two players' posteriors will have differing supports so long as the equilibrium beliefs in the original game are not additive.

The question of what epistemic condition would justify an equilibrium with non-additive beliefs has also been addressed in Epstein and Wang (1996) and Epstein (1997). What is different with the present paper is that we approach the question *on the presumption* that a non-additive measure (more accurately a non-additive belief function; see below) is but a reduced form representation of an additive belief on a state space that is not completely described in the given model. We take a different tack, in the sense that we focus on the implicit additive beliefs. From these "revealed beliefs", so to speak, we then reconstruct the degree of common belief inherent in the game. The cited literature and the present paper, essentially, work out the answers that two different ways of understanding a non-additive belief give, to the question of epistemics of equilibrium, in games with non-additive beliefs. In the former a non-additive belief is (part of) an irreducible description of preferences, as perhaps is the orthodox view; whereas, this paper seeks to answer the question for the case where one takes seriously the alternative view that it is a reduced form of what is actually a standard belief.

Defining a strategic equilibrium under Knightian uncertainty involves making several non-trivial modeling choices - namely, what specific class of measures to use to represent beliefs, whether to allow for a strict preference for randomization, whether to fix actions explicitly in the description of the equilibrium, or whether, instead of explicitly describing actions, to simply describe the supports of the beliefs, and if it is the latter, to further choose among the various possible notions of a support (see Ryan (1999) for a perspective on this choice). Unsurprisingly, the definition of equilibrium varies across the literature, each definition involving a particular set of modeling choices. The results in this paper are restricted to apply to a class of non-additive beliefs called belief functions, a restriction that is stronger than is standard in the literature on games with non-additive beliefs. We explain the reasons for the restriction in the next section. However, we do consider different definitions of equilibrium, covering the range of modeling choices mentioned above.

The rest of the paper is organized as follows. The next section briefly introduces the Choquet expected utility (CEU) model, the missing-states interpretation of CEU and the issue of CEU evaluation of randomized strategies. Section 3 analyses a definition of equilibrium with non-additive beliefs that allows for a strict preference for randomization while Section 4 analyses definitions of equilibrium that do not allow for such a strict preference. Section 5 concludes the paper with some remarks comparing the lessons from the analyses of the various definitions of equilibrium.

2 Preliminaries

2.1 Non-additive beliefs and incomplete state spaces

A non-additive measure μ defined on a sample space S is said to be a *belief function* if $\mu(S) = 1$, $\mu(\emptyset) = 0$ and there is a function $\alpha : 2^S \rightarrow [0, 1]$ such that, for any event $A \subseteq S$,

$$\mu(A) = \sum_{B \subseteq A} \alpha(B). \quad (1)$$

Our analysis in this paper is concerned with such belief functions. The interpretation of $\alpha(B)$, discussed by Dempster (1967), and popularized by Shafer (1976), is that of the “direct” evidence in support of the event B which cannot be further specified in terms of finer distinctions of B . Then, the overall belief put on the event A is given by the sum of all the direct support for propositions which imply A . Indeed, for any belief function μ , there is a unique non-negative α which satisfies (1), since α can be derived from μ as

$$\alpha(A) = \sum_{B \subseteq A} (-1)^{|A|-|B|} \mu(B). \quad (2)$$

Shafer (1976) refers to α as the *basic probability assignment* associated with μ . We will follow this terminology here.

Schmeidler (1989) pioneered the use of non-additive measures in decision theory to model uncertainty (or ambiguity) aversion. The force of the applicability of non-additive measures to decision theory lies in the way uncertainty aversion is implicit in the expression for the expected payoff defined as the *Choquet integral* $\int f d\mu$ with respect to a belief function μ , where f is a payoff function defined on the sample space S , and the integral $\int f d\mu$ is defined as follows (Choquet 1954):

$$\int f d\mu \equiv \int_0^\infty \mu(\{s | f(s) \geq t\}) dt + \int_{-\infty}^0 [\mu(\{s | f(s) \geq t\}) - \mu(S)] dt. \quad (3)$$

As was mentioned in the introduction, the understanding that non-additive expected utility is founded in situations where decision makers believe that they know the relevant state space only incompletely, inspires much of the analysis in this paper. Next, we recall the key insights of this incomplete state space interpretation of non-additive expected utility. (See Ghirardato (2001) and Mukerji (1997) for details.)

In the Savage set-up, a “state of the world” is a truth assignment to an *exhaustive* list of propositions describing aspects of the world to be. A way to capture the epistemic conditions that give rise to non-additive belief functions would be to depart from the Savage set-up to a world where DM *knows* he is able to take into account only a *partial* list of the relevant propositions. The DM is aware that the finest descriptions of contingent events he can think of, his ‘subjective’ states, are merely cells of a partition of the true ‘objective’ state space. Since a state may not be an exhaustive description, as the DM knows it, an act may at best map a state to a *set* of outcomes. Thus a prior on the states induces a probability function on the power set of the outcome space. A belief function is an equivalent representation of this induced probability. If $\alpha(A)$ is the basic probability assignment on A , a subset of the outcome space, then $\alpha(A)$ is the prior on the (subjective) state of the world (call it $\omega(A)$) corresponding to which the DM

can narrow down outcomes to the set A . Even though he suspects that he has failed to put some relevant aspects on his ‘list’ formulating $\omega(A)$, he does not know what those missing items may be. Hence, he cannot intelligently ‘distribute’ $\alpha(A)$ among the different outcomes in A . A striking result⁴ in Gilboa and Schmeidler (1994, theorem 4.3) presented as Lemma 1, below, shows, the Choquet evaluation of an act simply corresponds to taking a standard expectation combined with the ‘conservative’ procedure of associating $\alpha(A)$ with that outcome in A which yields least outcome from the act.

Lemma 1 $\int f d\mu = \sum_{A \subseteq S} \alpha(A) \min_{s \in A} f(s)$.

Actually, even if μ were a more general non-additive measure such as a *capacity* satisfying $\mu(S) = 1$, $\mu(\emptyset) = 0$ and $S \supseteq A \supseteq B \Rightarrow \mu(A) \geq \mu(B)$, the representation obtained in Lemma 1 still holds. However, the α ’s obtained from the decomposition of μ using (2) need not be non-negative if μ were a more general non-additive measure than a belief function. Hence, it is only when μ is restricted to belief functions that the corresponding α have a interpretation as a probability measure. Therefore it is only in the case of belief functions that Lemma 1 provides a basis for interpreting the Choquet evaluation an act w.r.t. μ as a standard expectation, taken with respect to the additive prior α , after extending the definition of the act to an augmented set of contingencies. The method we apply in this paper to turn a complete information game with non-additive beliefs into an incomplete information game with additive beliefs rests crucially on the device of decomposing a non-additive measure using (2) and applying Lemma 1; for the method to work the non-additive beliefs have to be belief functions, since for more a general measure the decomposition will not necessarily yield a probability measure.⁵

2.2 The issue of preference for randomization

Let \mathcal{F} be a set of acts, with f , a payoff function defined on the sample space S , being a typical element. Let $\Delta(\mathcal{F})$ be the set of all probability distributions over \mathcal{F} . Consider a DM evaluating a randomized strategy, say the mixed strategy σ which is a probability distribution over the elements of \mathcal{F} . A belief function μ describes the DM’s belief over S . It is a matter of some debate in the decision theory literature as to how a DM’s evaluation of a mixed strategy should be modelled, given that his belief on S is non-additive. Ghirardato and Le Breton (1998) provides

⁴The result has been noted by a number of other authors such as Shafer (1981), Jaffray (1989), Wasserman (1990), Smets (1981) and Dubois-Prade (1985). To get an intuition for the result notice that a belief function is a unique convex combination of unanimity games, and the Choquet integral of a function w.r.t. the unanimity game on A is the min of the function over A .

⁵Could our results be extended to more general non-additive beliefs? Ghirardato and Le Breton (2000) present a method of decomposing general capacities into additive probabilities. They then apply this to the context of a strategic form game to show that a strategy is a best response to a general capacity if and only if it is a best response to a belief function. However, this implication is of limited use in the context of an equilibrium. This is because two different acts which are best responses to the same capacity may not be best responses to the same belief function. However, the decomposition result itself holds more substantive ground for optimism. There are hurdles to overcome, however: one, the decomposition is not unique, since more than one additive prior may represent a given capacity, and two, this decomposition of general capacities would lead to a Bayesian game that was far more complicated than the one corresponding to belief functions. For more on the latter point, see footnote 10.

a very insightful account of this debate. The present discussion is an abbreviated version of that account.

The framework in Schmeidler (1989) assumes that the DM will evaluate $\sigma \in \Delta(\mathcal{F})$ according to the functional

$$\int_S \left[\sum_{f \in \mathcal{F}} \sigma(f) f(s) \right] d\mu. \quad (4)$$

In contrast to the standard additive case, though, the order of integration matters for Choquet integrals. Hence (4) is in general not the same as the functional which reverses the order of integration

$$\sum_{f \in \mathcal{F}} \sigma(f) \left[\int_S f(s) d\mu \right]. \quad (5)$$

Expression (4) evaluates mixed strategies as if the state s is decided first, and then the randomization is performed. Expression (5) does the opposite. Preference for randomization occurs when the DM evaluates in accordance with (4), it does not if (5) is used instead. It is not obvious as to which method of evaluation takes conceptual precedence: for instance, while Eichberger and Kelsey (1996) argue in favor of (5), Klibanoff (2001) shows that an intuitive modeling of the independence between σ and μ makes a strong case for using (4). Unsurprisingly, therefore, the literature applying non-additive beliefs to game theory follows a dual strand: one strand opting for evaluation in accordance with (4) while the other follows (5). The analysis in this paper is applied to both strands.

3 First Definition of Equilibrium

This section analyses a notion of equilibrium with non-additive beliefs which embodies preference for randomization. The equilibrium notion analyzed, is inspired by a notion introduced in Klibanoff (1996; page 8). Klibanoff used the maxmin multiple prior framework (Gilboa-Schmeidler 1989) in his work whereas this paper applies the Choquet expected utility model, using belief functions to model individuals' beliefs. At a minimum, the definition of equilibrium presented below is a translation of Klibanoff's definition; but, actually it goes beyond that. More specifically, if we were to replace each belief function in the definition below with the set of probabilities in the core of that belief function⁶ we would obtain a (slight) refinement of Klibanoff's version. The refinement is due to the third condition invoked in the definition below, which does not have a counterpart in Klibanoff's original notion. As will be explained in more detail in due course, this refinement facilitates a more meaningful interpretation of the equilibrium *in terms* of the incomplete information game we construct in our analysis.

For a two player finite game $G = (S_1, S_2, u_1, u_2)$, where S_i is player i 's strategy set and u_i is i 's payoff function, consider a belief function μ_1 defined on S_1 and μ_2 defined on S_2 . μ_j describes a player i 's belief about j 's choice of pure strategy. Let M_i be the set of all mixed strategies with support in S_i .

⁶The *core* of a belief function μ , (notation: $\Delta(\Omega)$ is the set of all additive probability measures on Ω) is defined as $\mathcal{C}(\mu) = \{m \in \Delta(\Omega) \mid m(E) \geq \mu(E), \text{ for all } E \subset \Omega\}$. Hence, $\mu(E) = \min_{m \in \mathcal{C}(\mu)} m(E)$. It may be useful to observe that the Choquet preferences used in this section fall in the intersection of the Choquet and maxmin models.

Definition 1 The quadruple $(\mu_1, \mu_2; \sigma_1, \sigma_2)$, where (μ_1, μ_2) is a pair of belief functions and (σ_1, σ_2) a pair of mixed strategies, is an equilibrium of G under Knightian uncertainty if,

- σ_i maximizes player i 's expected payoff, compared to all other mixed strategies in M_i , when player i holds belief μ_j over player j 's pure strategies, $i, j = 1, 2$.
- $\sigma_i \in \{m_i \in M_i \mid m_i(X_i) \geq \mu_i(X_i), \forall X_i \subseteq S_i\}$, $i = 1, 2$.
- $\sigma_i(s_i) > 0 \Rightarrow \mu_i(s_i) > 0, \forall s_i \in S_i, i = 1, 2$.

The expected payoff to player 1 of the mixed strategy σ_1 given beliefs μ_2 is the Choquet integral of the payoff function $\sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, \cdot)$, considered as a random variable over player 2's strategy set S_2 , with respect to the belief function μ_2 . That is, $\int_{S_2} [\sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, \cdot)] d\mu_2$, in accordance with the functional as described in (4). Analogous expressions hold for player 2's expected payoff.

The equilibrium is thus defined by a quadruple, elements of which are mutually related by three consistency requirements. The two classes of elements in the quadruple are equilibrium beliefs of each player about the other player's choice of strategy (μ_1, μ_2) and the actual strategies chosen by each player (σ_1, σ_2) . The first consistency requirement is that a player i 's strategy σ_i is a best response to his belief μ_j , where the evaluation of best response admits a strict preference for randomization as in (4). The final two consistency requirements are inspired by the notion that equilibrium beliefs must not be contradicted by actual equilibrium behavior. Hence, σ_i must be in the set of probabilities (or, mixed strategies) consistent with the belief μ_j . And, if a player plays a particular pure strategy with positive probability in equilibrium then his opponent's equilibrium belief function must assign a positive weight on that pure strategy.

3.1 Example I

As has been already noted our analysis and interpretation is driven by a construction of a game of incomplete information, where players have standard Bayesian beliefs, which mimics the original strategic situation (i.e., as described in the equilibrium with Knightian uncertainty) in its essential features. The construct will be called an *induced Bayesian game*. We now present an example of such a construction to illustrate the main ideas.

Our example is based on the following 3×2 game of complete information.

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

In this game, per “standard” analysis, R is dominated by L , and B is dominated by the mixed strategy which randomizes equally between T and M . Indeed, it has a standard solution in terms of two rounds of deletion of strictly dominated strategies which picks out (T, L) . However, there is an equilibrium of this game under Knightian uncertainty (per definition 1) in which player 1

[illegible]

There is common a -belief that the original game is being played, a being the minimum of player 1's posterior (a), and player 2's posterior (1).

At state $(0, 0)$, the strategy spaces and the payoffs of the two players are given by the original game $G = (S_1, S_2, u_1, u_2)$. At state $(0, Q)$, $Q \subseteq S_2$, an extensive form game is played in which player 1 moves first, and player 2 moves after observing player 1's chosen random strategy σ_1 *but not the outcome of the randomization*. Player 2's action set is given by Q . The payoff function of player 1 is identical to the payoff function u_1 in G , except that it is restricted to $S_1 \times Q$. The payoff function of player 2, also defined on $S_1 \times Q$ is simply the negative of player 1's payoff. At state $(Q, 0)$, $Q \subseteq S_1$, an extensive form game is played which is analogous to the extensive form game at game $(0, Q)$, except that the roles of the two players are reversed. The example here is driven by the game at state $(0, \{L, R\})$. This is the game in which player 1 moves first, and player 2 moves after observing 1's choice of (randomized) strategy (though not the outcome of the randomization), and where the payoffs are given by

		2	
		L	R
1	T	3, -3	0, 0
	M	0, 0	3, -3
	B	1, -1	1, -1

By our construction, in the state $(0, \{L, R\})$, if T is played by 1, player 2 responds with R , while if M is played, 2 responds with L . If either B or σ_1 is played, 2 is indifferent between L and R . (In the case 1 plays σ_1 , 2 expects a payoff of $-3/2$ whatever he chooses.) Thus, solving backwards, 1's payoff from playing T or M is zero, while he can guarantee a payoff of 1 if he plays B and $3/2$ if he plays σ_1 . Provided that the probability of the state $(0, \{L, R\})$ is large enough (i.e., if $1 - a > 1/2$), player 1 finds it *strictly* optimal to play σ_1 .

It is instructive to note how the preference for randomization comes through in the construction of the induced Bayesian game. In the induced Bayesian game, the possible types of player 2 do not get to see the outcome of the randomization before they choose their action. Given player 1's "paranoia" that each crazy type of player 2 will inflict the worst damage it can, randomization serves to disguise his play and hence foil the player 2's attempt to do the worst. For instance in the example, mixing allows player 1 to limit his (expected) damages to 1.5. On the hand, if the information structure of the induced Bayesian game was such that the crazy types (i.e., the non-zero types) chose with full knowledge of the outcome of the randomization resorted to by their opponents, then randomization would not serve as a "veil" and there would be, unsurprisingly, no strict preference for randomization. Indeed, this is the key difference in the construction of the induced Bayesian game corresponding to notions of equilibrium which do not admit preference for randomization (e.g., Dow and Werlang (1994)).

3.2 Induced Bayesian Game I

Having dealt with the preliminaries, we are now ready to present our main construction. We present an algorithm for constructing an incomplete information game with standard *additive* beliefs for any given situation of Knightian uncertainty, and demonstrate how equilibrium under Knightian uncertainty has an interpretation in terms of an equilibrium in this standard framework.

Let the underlying two player, finite normal form game be $G = (S_1, S_2, u_1, u_2)$. Consider the quadruple $(\mu_1, \mu_2; \sigma_1, \sigma_2)$, where (μ_1, μ_2) are belief functions and (σ_1, σ_2) a pair of mixed strategies which satisfy the following consistency conditions (a) and (b) which, we recall, were also incorporated in the definition of an equilibrium under Knightian uncertainty:

- (a) $\sigma_i \in \{m_i \in M_i \mid m_i(X_i) \geq \mu_i(X_i), \forall X_i \subseteq S_i\}, i = 1, 2.$
- (b) $\sigma_i(s_i) > 0 \Rightarrow \mu_i(s_i) > 0, \forall s_i \in S_i, i = 1, 2.$

We will construct a Bayesian game as follows. We begin with the type space.

- (1) Type Space.** Denote by \mathcal{Q}_i the collection of all non-empty subsets of S_i . The type space of player i is defined to be the set

$$T_i \equiv \mathcal{Q}_i \cup \{0\}.$$

The *type space* T for the Bayesian game is the set of all pairs (t_1, t_2) such that both of the following hold.

- $t_1 \in T_1$ and $t_2 \in T_2$,
- Either $t_1 = 0$ or $t_2 = 0$.

Each player knows his own type (and only his own type). Hence, player 1 can distinguish between states (t_1, t_2) and (t'_1, t'_2) if and only if $t_1 \neq t'_1$. Similarly, player 2 can distinguish between states (t_1, t_2) and (t'_1, t'_2) if and only if $t_2 \neq t'_2$.

- (2) Common Prior.** There is a common prior p over the type space T . The probability of the state (t_1, t_2) is denoted by $p(t_1, t_2)$, and is defined as follows. Define the number

$$\eta_i = \min_{s_i \in \text{support}(\sigma_i)} \left\{ \frac{\mu_i(\{s_i\})}{\sigma_i(s_i)} \right\} \quad (6)$$

Note, our assumptions imply that $0 < \eta_i \leq 1$. The consistency condition (a) implies $\eta_i \leq 1$. Consistency condition (b) implies $\eta_i > 0$. We denote by α_i the basic probability assignment implied by the belief function μ_i as shown in (2). Then $p(t_1, t_2)$ is defined as follows.

- $p(0, 0) = \frac{\eta_1 \eta_2}{\eta_1 + \eta_2(1 - \eta_1)}$
- $p(0, Q) = \begin{cases} \frac{\eta_1 [\mu_2(\{s_2\}) - \eta_2 \sigma_2(s_2)]}{\eta_1 + \eta_2(1 - \eta_1)} & \text{if } Q \text{ is the singleton } \{s_2\} \\ \frac{\eta_1 \alpha_2(Q)}{\eta_1 + \eta_2(1 - \eta_1)} & \text{if } |Q| \geq 2, Q \in \mathcal{Q}_2 \end{cases}$
- $p(Q, 0) = \begin{cases} \frac{\eta_2 [\mu_1(\{s_1\}) - \eta_1 \sigma_1(s_1)]}{\eta_1 + \eta_2(1 - \eta_1)} & \text{if } Q \text{ is the singleton } \{s_1\} \\ \frac{\eta_2 \alpha_1(Q)}{\eta_1 + \eta_2(1 - \eta_1)} & \text{if } |Q| \geq 2, Q \in \mathcal{Q}_1 \end{cases}$

We first check that these expressions are non-negative and sum to 1. Note, since $0 < \eta_i \leq 1$, it follows that $\eta_1 + \eta_2(1 - \eta_1) > 0$. Since η_1 and η_2 (and α_i , since μ_i is a belief function), are non-negative $p(t_1, t_2)$ is non-negative for all pairs of types. To show that they sum to 1, denote by \mathcal{Q}_i^* the collection of subsets of S_i that have two or more elements. Then,

$$\begin{aligned}
\sum_{t_2 \in T_2} p(0, t_2) &= \frac{\eta_1}{\eta_1 + \eta_2(1 - \eta_1)} \left\{ \eta_2 + \sum_{s_2 \in S_2} (\mu_2(\{s_2\}) - \eta_2 \sigma_2(s_2)) + \sum_{Q \in \mathcal{Q}_2^*} \alpha_2(Q) \right\} \\
&= \frac{\eta_1}{\eta_1 + \eta_2(1 - \eta_1)} \left\{ \eta_2 + \sum_{s_2 \in S_2} \mu_2(\{s_2\}) - \eta_2 \overbrace{\sum_{s_2 \in S_2} \sigma_2(s_2)}^{=1} + \sum_{Q \in \mathcal{Q}_2^*} \alpha_2(Q) \right\} \\
&= \frac{\eta_1}{\eta_1 + \eta_2(1 - \eta_1)} \left\{ \sum_{s_2 \in S_2} \mu_2(\{s_2\}) + \sum_{Q \in \mathcal{Q}_2^*} \alpha_2(Q) \right\} \\
&= \frac{\eta_1}{\eta_1 + \eta_2(1 - \eta_1)} \sum_{A \subseteq S_2} \alpha_2(A) \\
&= \frac{\eta_1}{\eta_1 + \eta_2(1 - \eta_1)}
\end{aligned}$$

where we have used the fact that for a singleton event $\{s\}$, $\mu_i(\{s\}) = \alpha_i(\{s\})$, and $\sum_{A \subseteq S_2} \alpha_2(A) = \mu_2(S_2) = 1$ by (1). Similarly, we have

$$\sum_{t_1 \in T_1} p(t_1, 0) = \frac{\eta_2}{\eta_1 + \eta_2(1 - \eta_1)},$$

so that

$$\sum_{t \in T} p(t) = \sum_{t_2 \in T_2} p(0, t_2) + \sum_{t_1 \in T_1} p(t_1, 0) - p(0, 0) = 1.$$

Hence p is a probability distribution over the type space T .

(3) Payoffs. We complete the description of the Bayesian game by defining the payoffs at each state $t \in T$.

- At $t = (0, 0)$, the strategy spaces and the payoffs of the two players are given by the original game $G = (S_1, S_2, u_1, u_2)$.
- At $t = (0, Q)$, an extensive form game is played in which player 1 moves first, and player 2 moves after observing player 1's chosen random strategy σ_1 *but not the outcome* of the randomization. Player 2's action set is given by Q . The payoff function of player 1 is identical to the payoff function u_1 in G , except that it is restricted to $S_1 \times Q$. The payoff function of player 2, also defined on $S_1 \times Q$ is given by

$$-u_1. \tag{7}$$

In other words, player 2's payoff is the negative of player 1's payoff.

- At $t = (Q, 0)$, an extensive form game is played, which is analogous to the game above except that the roles of the two players are reversed. It is player 2 who now moves first and player 1 moves second, after observing player 2's chosen random strategy σ_2 *but not the outcome* of the randomization. Player 2's action set is S_2 , identical to game G . Player 1's action set is given by Q . The payoff function of player 2 is identical to the payoff function u_2 in G except that it is restricted to $Q \times S_2$. The payoff function of player 1, is given by $-u_2$.

This exhausts all the possible states in the type space. Note that the action sets at each state are such that a player cannot infer anything more about the underlying state from observing his own action set, beyond that given by the knowledge of his own type. For player 1, his action set is given by S_1 whenever he is of type 0. No information is revealed about the type of player 2. Similar comments apply to player 2. Bringing all the strands of the construction together, we say that the Bayesian game obtained from the above construction is the *induced Bayesian game*, induced by $(\mu_1, \mu_2; \sigma_1, \sigma_2)$ and G .

First Representation Theorem *Fix a quadruple $(\mu_1, \mu_2; \sigma_1, \sigma_2)$ such that the mixed strategy σ_i is in the core of the belief function μ_i and such that a positive weight in the mixed strategy σ_i implies a positive weight in the belief function μ_i . Then the quadruple is an equilibrium of G under Knightian uncertainty according to definition 1 if and only if the mixed strategies (σ_1, σ_2) are played at state $(0, 0)$ in an equilibrium of the induced Bayesian game.*

Proof: We will take the proof of this result in stages. Let us first consider the optimal action of type 0 of player i in the induced Bayesian game. For concreteness, we will put ourselves in the shoes of type 0 of player 1. Consider the game at state $(0, Q)$, where $Q \neq 0$. From our construction of the induced game, player 2 moves after observing player 1's choice of randomized strategy though not the realization of the randomization. Player 2's action set is given by $Q \subseteq S_2$, and 2's payoff function is given by $-u_1$. Thus, if player 1 chooses σ_1 , 1's payoff at the state $(0, Q)$ is

$$\min_{s_2 \in Q} \sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, s_2).$$

At state $(0, 0)$, the normal form game G is played. If player 2 plays the mixed strategy σ_2 , then the expected payoff to player 1 of playing σ_1 is given by

$$\sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \sigma_1(s_1) \sigma_2(s_2) u_1(s_1, s_2).$$

When Q has two more elements, the probability of state $(0, Q)$ conditional on player 1 being type 0 is given by

$$\frac{p(0, Q)}{\sum_{t_2 \in T_2} p(0, t_2)} = \alpha_2(Q),$$

while if Q were a singleton, say $\{s_2\}$, then the conditional probability of $(0, \{s_2\})$, conditional on player 1 being type 0 is given by

$$\mu_2(\{s_2\}) - \eta_2 \sigma_2(s_2) \tag{8}$$

while the conditional probability of state $(0, 0)$ is

$$\eta_2 = \min_{s_2 \in \text{support}(\sigma_2)} \left\{ \frac{\mu_2(\{s_2\})}{\sigma_2(s_2)} \right\}$$

Hence, the expected payoff of type 0 of player 1 choosing to play a strategy σ_1 when type 0 of player 2 plays the mixed strategy σ_2 in state $(0, 0)$ is given by

$$\begin{aligned} & \eta_2 \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \sigma_1(s_1) \sigma_2(s_2) u_1(s_1, s_2) \\ & + \sum_{s_1 \in S_1} \sigma_1(s_1) \sum_{s_2 \in S_2} [\mu_2(\{s_2\}) - \eta_2 \sigma_2(s_2)] u_1(s_1, s_2) \\ & + \sum_{Q \in \mathcal{Q}_2^*} \alpha_2(Q) \min_{s_2 \in Q} \sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, s_2) \\ = & \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \sigma_1(s_1) \mu_2(\{s_2\}) u_1(s_1, s_2) \\ & + \sum_{Q \in \mathcal{Q}_2^*} \alpha_2(Q) \min_{s_2 \in Q} \sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, s_2) \end{aligned} \quad (9)$$

Since $\mu_2(\{s_2\}) = \alpha_2(\{s_2\})$ we may rewrite the r.h.s. of (9) more compactly as

$$\sum_{Q \in \mathcal{Q}_2} \alpha_2(Q) \min_{s_2 \in Q} \sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, s_2) \quad (10)$$

Thus, in order to show that an action σ_1 is optimal for type 0 of player 1 when type 0 of player 2 plays the induced mixed strategy σ_2 , we must show that σ_1 maximizes (10).

Meanwhile, from the hypothesis that $(\mu_1, \mu_2; \sigma_1, \sigma_2)$ is an equilibrium of G under Knightian uncertainty, we know that σ_1 maximizes expected payoff for player 1. That is, σ_1 maximizes

$$\int_{S_2} \left[\sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, \cdot) \right] d\mu_2. \quad (11)$$

Hence, if we could show that (11) is equal to (10), then our result would be proved. This is indeed the case. The equivalence of (11) and (10) follows from Lemma 1, which written in terms of strategies and utility payoffs would read,

$$\int u_1(s_1, \cdot) d\mu_2 = \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s_1, s_2). \quad (12)$$

Define $\tilde{u}(s_1, \cdot) \equiv \sum_{s_1 \in S_1} \sigma_1(s_1) u_1(s_1, \cdot)$ and notice we may therefore rewrite (11) as $\int \tilde{u}(s_1, \cdot) d\mu_2$ and (10) as $\sum_{Q \in \mathcal{Q}_2} \alpha_2(Q) \min_{s_2 \in Q} \tilde{u}(s_1, s_2)$. Now, apply (12) to $\int \tilde{u}(s_1, \cdot) d\mu_2$ to obtain the required equivalence of (11) and (10). This completes the proof of the theorem. ■

The first representation theorem shows that an equilibrium under Knightian uncertainty may be represented as a strategic equilibrium where players are Bayesian, i.e., they have probabilistic beliefs and are responding to a situation of incomplete information. In this representation, as player i sees it, there are several contingent states of the world. While the game played in state $(0, 0)$ is as described by G with his opponent playing the strategy σ_j , player i thinks that

different games are possible at states $(0, t_j)$. The other possibilities are similar to assuming that an opponent may be one of several “crazy types”, a standard trick in economic modelling. However, while a “standard” crazy type is typically committed to a particular pure strategy, here a crazy type may only be pinned down to a particular *set* of strategies. Further, the crazy type t_j is assumed to know player i ’s randomized strategy and, given this information, chooses that strategy (from the set Q , the set of strategies in S_j associated with that state) which minimizes player i ’s payoff. Finally, i ’s posterior on the state $(0, t_j)$ is the basic probability assignment on Q , as evident in μ_j , i ’s equilibrium belief as found in the Knightian equilibrium of the original game. Thus, i ’s posterior distribution over S_j in the equilibrium of the Bayesian game is not σ_j ; it is a belief which mimics μ_j in its essentials and which while allowing for σ_j allows for other possibilities too – a perturbation of σ_j as it were. An interpretation of i ’s belief μ_j as may be evinced from the induced Bayesian game, could be that i is not sure of j ’s motives or even the (informational and sequential) structure of the game and is inclined to take a particularly pessimistic view of the possibilities, including the possibility that j is actually out to get him.

3.3 Common Belief of Game G

One of the consequences of our representation theorem is that the equilibrium beliefs and actions embodied in $(\mu_1, \mu_2; \sigma_1, \sigma_2)$ can be translated into orthodox game-theoretic notions as the beliefs and actions of a pair of *types* in a Bayesian game. By its very nature, the beliefs of a type in a Bayesian game are posterior beliefs conditional on knowing one’s own identity. If there were common knowledge of the types in a Bayesian game, there would be common knowledge of the payoffs, and we would be in the realms of classical game theory in which payoffs are common knowledge. In general, however, the updated beliefs of types at a given state have different supports. One of the types will allow some states to be possible which the other type rules out, and conversely, the first type may rule out some states which the second type allows as possible.

Under these circumstances, propositions which are true at a given state need not be common knowledge at that state. In terms of the Bayesian game induced by $(\mu_1, \mu_2; \sigma_1, \sigma_2)$, even if the underlying game being played is the benchmark game G , it will not be common knowledge that this is the case, in general. The deviation from common knowledge will depend on the deviation of the belief function(s) from additivity. There will there be common knowledge of the game being played if and only if the beliefs functions are conventional, *additive* probability measures.

To state this claim more precisely, we appeal to the notion of common p -belief. Say that an individual p -believes an event A if this individual puts probability at least p to this event. An event A is said to be *common p -belief* among a set of individuals if everyone p -believes A , everyone p -believes that everyone p -believes A , and so on for every order of belief. We have the following feature of the Bayesian game induced by $(\mu_1, \mu_2; \sigma_1, \sigma_2)$.

Proposition 1 *Suppose $(\mu_1, \mu_2; \sigma_1, \sigma_2)$ is an equilibrium of G under Knightian uncertainty according to definition 1. Then, at state $(0, 0)$ of the induced Bayesian game, there is common λ -belief that the game G is being played, where*

$$\lambda = \min \{ \eta_1, \eta_2 \} .$$

For any $\lambda' > \lambda$, it is not common λ' -belief that the game G is being played.

If and only if at least one of the belief functions (μ_1, μ_2) is non-additive, we have $\lambda < 1$, so that there is a failure of common 1-belief of the game being played.⁷ In general, the lower is the weight given to the singleton sets $\{s_i\}$ by these belief functions, the greater will be the departure from common belief. Since our first representation theorem has identified the beliefs and actions embodied in $(\mu_1, \mu_2; \sigma_1, \sigma_2)$ with the actions and beliefs at the state $(0, 0)$, this result highlights the failure of common knowledge implicit in the notion of equilibrium under Knightian uncertainty, under our interpretation of the induced Bayesian game.

Since the induced Bayesian game has a common prior over the type space, the proof of the above result is relatively straightforward. We appeal to the following characterization of common p -belief given by Monderer and Samet (1989 p.177). An event A is said to be p -evident to a set of individuals if, whenever A is true, everyone puts probability at least p to A . Common belief has the following “fixed point” characterization in terms of p -evident events.

Lemma 2. An event A is common p -belief at the state t if and only if there is a p -evident event $E \subseteq A$ such that $t \in E$.

The proof of our theorem on the failure of common belief rests on the observation that the singleton event $\{(0, 0)\}$ is λ -evident at the state $(0, 0)$. At the state $(0, 0)$, a player j in the Bayesian game places posterior probability η_i to the event $\{(0, 0)\}$. Thus, for λ defined as the smaller of η_1 and η_2 , the event $\{(0, 0)\}$ is λ -evident at $(0, 0)$. Since the true payoffs are given by G on the event $\{(0, 0)\}$, we thus have common λ -belief at the state $(0, 0)$ that G is being played. For any value $\lambda' > \lambda$, it is not common λ' -belief that G is being played.

The following is an intuition for thinking of η_i as a measure of j 's belief about the likelihood of payoff functions and the information structure of the game being the same as in G . The ratio η_i gives the (minimum) relative weight that j 's belief function puts on s_i as compared to the weight put on it by the “rational” incarnation of player i . Hence, the higher is this relative weight, the more prominently the rational player's strategy figures in anchoring the belief (on singleton strategies). This intuition also shows that if the third consistency condition in the definition of equilibrium did not hold, it would not be meaningful to interpret the equilibrium as one where a player has incomplete information about the rationality of his opponent; without the condition, η_i could be zero, in which case it would be impossible for j to believe that i could be rational.

4 Second Definition of Equilibrium

Dow and Werlang's (1994) definition of equilibrium with non-additive beliefs is different from the one discussed in the last section in two key respects. One, it embodies no preference for randomization, and two, the equilibrium is essentially defined by a pair of beliefs. Behavior, that is, strategies are not explicitly described (or even pinned down implicitly). This latitude means that the representation in terms of the induced Bayesian game is not completely tied down without imposing further structure. In particular, much depends on the definition of the *support* of the belief functions. Indeed, the only restriction on equilibrium behavior is that only

⁷If μ_i is additive, then the core of μ_i coincides with μ_i ; hence, σ_i must coincide with μ_i , therefore implying that $\eta_i = 1$. Conversely, if $\eta_i = 1$, then $\mu_i(s_i) \geq \sigma_i(s_i)$. On the other hand, since σ_i is in the core of μ_i it must be that $\sigma_i(s_i) \geq \mu_i(s_i)$. Hence, $\eta_i = 1 \Rightarrow \sigma_i(s_i) = \mu_i(s_i)$, thereby implying that μ_i is additive.

those strategies which appear in the set defined to be the support of the equilibrium beliefs may be played in an equilibrium. There are at least three possible definitions⁸ of the support A of a belief function μ_i .

- A consists of all elements $s_i \in S_i$ such that $\mu_i(s_i) > 0$.
- A is the smallest event such that $\mu_i(A) = 1$.
- A is such that $\mu_i(A^c) = 0$ and $\mu_i(B^c) > 0$ for all $B \subset A$.

Let \tilde{X}_i denote the set of *all* strategies $s_i \in S_i$ such that $\mu_i(\{s_i\}) > 0$. If we used the first definition of support, then \tilde{X}_i would be the support of μ_i . In general, however, we allow the support X_i to be different from \tilde{X}_i (though, $\tilde{X}_i \subseteq X_i$). For any of the three definitions of support we can define equilibrium in a two player finite game $G = (S_1, S_2, u_1, u_2)$ as follows.

Definition 2. The pair of belief functions (μ_1, μ_2) with respective supports (X_1, X_2) is an equilibrium of G under Knightian uncertainty if,

- for every $s_1 \in X_1$, s_1 maximizes player 1's expected payoff, compared to all other pure strategies in S_1 , when player 1 holds belief μ_2 over player 2's strategies,
- for every $s_2 \in X_2$, s_2 maximizes player 2's expected payoff, compared to all other pure strategies in S_2 , when player 2 holds belief μ_1 over player 1's strategies.
- $\tilde{X}_i \neq \emptyset$.

The expected payoff to player 1 of action s_1 given beliefs μ_2 is the Choquet integral of the payoff function $u_1(s_1, \cdot)$ considered as a random variable over player 2's strategy set S_2 with respect to the belief function μ_2 . That is, $\int u_1(s_1, \cdot) d\mu_2$. Analogous expressions hold for player 2's expected payoff. Note, mixed strategies are evaluated as in (5).

Definition 2 differs from Dow and Werlang's definition of equilibrium by the inclusion of the proviso that $\tilde{X}_i \neq \emptyset$, saying that there is at least one strategy for each player which receives a positive weight as a singleton in the opponent's (equilibrium) belief function. Without the proviso it is not possible for us to adduce a meaningful interpretation to the induced Bayesian game we construct to represent this definition of equilibrium, since it plays a role in the interpretation analogous to the role played by the third consistency condition in definition 1.

4.1 Example II

We will re-work our previous example through the Dow-Werlang equilibrium notion to highlight some similarities and differences in the induced Bayesian game we will construct to analyze such equilibria, as compared to the construction corresponding to definition 1. We first recall the 3×2 game used earlier:

⁸These are the definitions found in the literature. The first appears in Marinacci (2000), the second is implicit in Ghirardato and Le Breton (1998), while the third was first proposed in Dow and Werlang (1994).

		2	
		<i>L</i>	<i>R</i>
1	<i>T</i>	3, 1	0, 0
	<i>M</i>	0, 1	3, 0
	<i>B</i>	1, 1	1, 0

Recall, the standard solution in terms of two rounds of deletion of strictly dominated strategies which picks out (T, L) . However, we now show that there is an equilibrium of this game under Knightian uncertainty (per definition 2) in which player 1 plays B , assuming we apply the third definition of support, following Dow-Werlang (1994). The equilibrium is given by the quadruple (μ_1, μ_2, X_1, X_2) , where $\mu_1(\{B\}) = 1$, $\mu_2(\{L\}) = a < 1/3$, $\mu_2(\{R\}) = 0$, $\mu_2(\{L, R\}) = 1$, and $X_1 = \{B\}$, $X_2 = \{L\}$.

For player 2, L is the optimal action given beliefs μ_1 , so that one half of the condition for equilibrium is satisfied. For player 1, the expected payoff to playing B is given by $\int u_1(B, \cdot) d\mu_2$, which lemma 1 tells us is equal to

$$a \cdot 1 + (1 - a) \cdot \min\{1, 1\} = 1.$$

Thus, action B yields an expected payoff of 1. This contrasts with the expected payoff from T , which is

$$a \cdot 3 + (1 - a) \cdot \min\{3, 0\} = 3a.$$

Since $a < 1/3$, the expected payoff from T falls short of playing B . Finally, playing M yields

$$a \cdot 0 + (1 - a) \cdot \min\{0, 3\} = 0,$$

which is clearly inferior to B . Hence, B maximizes 1's expected payoff in terms of the Choquet integral given beliefs μ_2 . Hence, the quadruple above constitutes an equilibrium of the game under Knightian uncertainty per definition 2.

To understand what's going on here, it is illuminating to see what drives the beliefs of the players in the induced Bayesian game at state $(0, 0)$. The type space of the induced Bayesian game is represented below; the reader will note the type space of the induced Bayesian game here does not include the singleton subsets of the strategy space.

$(0, 0)$	$(0, \{L, R\})$
$(\{T, M\}, 0)$	
$(\{T, B\}, 0)$	
$(\{M, B\}, 0)$	
$(\{T, M, B\}, 0)$	

Player 1 can distinguish between the rows of this array, while player 2 can distinguish between the columns. The common prior over this type space is given by

a	$1 - a$
0	
0	
0	
0	

There is common a -belief that the original game is being played, a being the minimum of player 1's posterior (a), and player 2's posterior (1). The example is driven by the game at state $(0, \{L, R\})$. This is the game in which player 1 moves first, and player 2 moves after observing 1's choice (including the actual outcome of a randomization), and where the payoffs are given by

		2	
		L	R
1	T	3, -3	0, 0
	M	0, 0	3, -3
	B	1, -1	1, -1

If T is played by 1, player 2 responds with R , while if M is played, 2 responds with L . If B is played, 2 is indifferent between L and R . Thus, solving backwards, 1's payoff from playing T or M is zero, while he can guarantee a payoff of 1 if he plays B . Provided that the probability of the state $(0, \{L, R\})$ is larger enough (if $1 - a > 2/3$), player 1 finds it optimal to play B .

This example shows how non-additive beliefs introduces an apparent violation of the key feature in standard game theory of “taking the other's action as given”. In this example, when player 1 contemplates a deviation from T to M , say, the opponent's action *cannot* be taken as given in this state. Since the force of the dominance principle lies in the fact that opponent's actions do not change with one's own deviation, it is perhaps not surprising that we can generate an example such as above. This of course is the consequence of player 1 playing as if a “crazy” type of player 2 not only employs a maxmin strategy but also actually gets to *know* the *outcome* of 1's randomization before making his own choice. Also, an immediate implication of the assumed information structure is that randomization no longer serves as a disguise and there is no more a reason to strictly prefer to randomize.

4.2 Induced Bayesian Game II

Let the underlying two player, finite normal form game be $G = (S_1, S_2, u_1, u_2)$. Consider the quadruple (μ_1, μ_2, X_1, X_2) , where μ_i is a belief function on S_i with support X_i such that $X_i \supseteq \tilde{X}_i$ and $\tilde{X}_i \neq \emptyset$. We will construct a Bayesian game as follows.

- (1) **Type Space.** Denote by \mathcal{Q}_1^* the collection of subsets of S_1 with two or more elements. Similarly, denote by \mathcal{Q}_2^* the collection of subsets of S_2 with two or more elements. The type space of player i is defined to be the set

$$T_i \equiv \mathcal{Q}_i^* \cup \{0\}.$$

The *type space* T for the Bayesian game is the set of all pairs (t_1, t_2) such that both of the following hold.

- $t_1 \in T_1$ and $t_2 \in T_2$,
- Either $t_1 = 0$ or $t_2 = 0$.

Each player knows his own type (and only his own type). Hence, player 1 can distinguish between states (t_1, t_2) and (t'_1, t'_2) if and only if $t_1 \neq t'_1$. Similarly, player 2 can distinguish between states (t_1, t_2) and (t'_1, t'_2) if and only if $t_2 \neq t'_2$.

(2) Common Prior. The common prior p over the type space T is defined as follows. Use the shorthand notation λ_i for:

$$\lambda_i \equiv \sum_{s_i \in S_i} \mu_i(\{s_i\}). \quad (13)$$

Then,

- $p(0, 0) = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2(1 - \lambda_1)}$
- $p(0, Q) = \frac{\lambda_1 \alpha_2(Q)}{\lambda_1 + \lambda_2(1 - \lambda_1)}, Q \in \mathcal{Q}_2^*$
- $p(Q, 0) = \frac{\lambda_2 \alpha_1(Q)}{\lambda_1 + \lambda_2(1 - \lambda_1)}, Q \in \mathcal{Q}_1^*$

The fact that μ_i are belief functions and that $\tilde{X}_i \neq \emptyset$ ensures that $0 < \lambda_i \leq 1$. Hence, these expressions are non-negative. To see that they sum to 1, note that

$$\begin{aligned} \sum_{t_2 \in T_2} p(0, t_2) &= \frac{\lambda_1 (\lambda_2 + \sum_{Q \in \mathcal{Q}} \alpha_2(Q))}{\lambda_1 + \lambda_2(1 - \lambda_1)} \\ &= \frac{\lambda_1 (\sum_{A \subseteq S_2} \alpha_2(A))}{\lambda_1 + \lambda_2(1 - \lambda_1)} \\ &= \frac{\lambda_1}{\lambda_1 + \lambda_2(1 - \lambda_1)}. \end{aligned}$$

where the final equality follows from the fact that $\sum_{A \subseteq S_2} \alpha_2(A) = \mu_2(S_2) = 1$ by (1). Similarly, we have

$$\sum_{t_1 \in T_1} p(t_1, 0) = \frac{\lambda_2}{\lambda_1 + \lambda_2(1 - \lambda_1)},$$

so that

$$\sum_{t \in T} p(t) = \sum_{t_2 \in T_2} p(0, t_2) + \sum_{t_1 \in T_1} p(t_1, 0) - p(0, 0) = 1.$$

Hence p is a probability distribution over the type space T .

(3) Payoffs. The payoffs of the game are identical to our previous definition except for one key feature, highlighted below.

- At $t = (0, 0)$, the strategy spaces and the payoffs of the two players are given by the original game $G = (S_1, S_2, u_1, u_2)$.

- At $t = (0, Q)$, $Q \in \mathcal{Q}_2^*$, an extensive form game is played in which player 1 moves first, and player 2 moves after observing player 1's action. If player 1 randomizes, then player 2 observes the realization of the randomization before choosing. Player 1's action set is S_1 , identical to the game G . Player 2's action set is given by Q . The payoff function of player 1 is identical to the payoff function u_1 in G , except that it is restricted to $S_1 \times Q$. The payoff function of player 2, also defined on $S_1 \times Q$ is given by

$$-u_1. \quad (14)$$

- At $t = (Q, 0)$, $Q \in \mathcal{Q}_1^*$, an extensive form game is played, which is analogous to the game above, except that the roles of the two players are reversed.

We say that the Bayesian game obtained from the above construction is the *induced Bayesian game*, induced by (μ_1, μ_2, X_1, X_2) and G .

Turning now to the actions implied by the belief functions, we make explicit the role of the belief functions as determining the actions of the players in the induced Bayesian game by defining a strategy called the *induced mixed strategy* for each player.

Definition 3 A probability distribution $\sigma_i(\delta)$ over S_i parameterized by $\delta \in [0, 1]$ is the *induced mixed strategy*, induced by (μ_1, μ_2, X_1, X_2) and G , if the probability weight given by $\sigma_i(\delta)$ to $s_i \in S_i$ is

$$\begin{cases} \frac{\mu_i(\{s_i\})}{\sum_{s'_i \in S_i} \mu_i(\{s'_i\})} & \text{if } s_i \in \tilde{X}_i \text{ and } X_i = \tilde{X}_i \\ \frac{\mu_i(\{s_i\})(1-\delta)}{\sum_{s'_i \in S_i} \mu_i(\{s'_i\})} & \text{if } s_i \in \tilde{X}_i \text{ and } X_i \neq \tilde{X}_i \\ \frac{\frac{\delta}{|X_i \setminus \tilde{X}_i|}}{\sum_{s'_i \in S_i} \mu_i(\{s'_i\})} & \text{if } s_i \in X_i \setminus \tilde{X}_i \\ 0 & \text{if } s_i \notin X_i \end{cases}$$

Definition 4 A *strong ϵ -equilibrium* of any strategic form game is a profile of mixed strategies, one for each player, such that a player's expected payoff from any pure strategy in the support of his mixed strategy in the profile is at most ϵ less than the expected payoff from any other strategy.

The above definition follows very much the spirit of the Radner (1980) definition of an ϵ -equilibrium, but is a stricter concept in the sense that while any *strong ϵ -equilibrium* will be a ϵ -equilibrium the converse is not true. The refinement follows from the fact that the above definition requires each pure strategy in the support of a mixed strategy to be *individually* an ϵ -best response.

We can now state our representation result for the Dow-Werlang version of an equilibrium as follows.

Second Representation Theorem Fix a quadruple (μ_1, μ_2, X_1, X_2) such that at least one strategy in the support of each belief function gets positive weight in the belief function, i.e., $\tilde{X}_i, \tilde{X}_i \subseteq X_i$, is non-empty. Then the quadruple is an equilibrium of G under Knightian uncertainty according to definition 2 if and only if there is a sequence $\{(\epsilon_n, \delta_n)\}$ with

$\epsilon_n \geq 0$, $\delta_n \geq 0$ and $\lim_{n \rightarrow \infty} (\epsilon_n, \delta_n) = (0, 0)$ such that, at every n , the induced mixed strategies $(\sigma_1(\delta_n), \sigma_2(\delta_n))$ are played at state $(0, 0)$ in a strong ϵ -equilibrium of the induced Bayesian game, where $\epsilon = \epsilon_n$. Further, if $X_i \neq \tilde{X}_i$ for some i , then $\epsilon_n > 0$, $\delta_n > 0$ for all n .

Proof. Let us first consider the optimal action of type 0 of player i in the induced Bayesian game. For concreteness, we will put ourselves in the shoes of type 0 of player 1. Consider the game at state $(0, Q)$, where $Q \in \mathcal{Q}_2^*$. From our construction of the induced game, player 2 moves after observing player 1's action, player 2's action set is given by $Q \subseteq S_2$, and 2's payoff function is given by $-u_1$. Thus, if player 1 chooses s_1 , his payoff at the state $(0, Q)$ is

$$\min_{s_2 \in Q} u_1(s_1, s_2).$$

At state $(0, 0)$, the normal form game G is played. If player 2 plays the induced mixed strategy $\sigma_2(\delta)$, then the expected payoff to player 1 of playing s_1 is given by

$$\sum_{s_2 \in S_2} \sigma_2(\delta)(s_2) u_1(s_1, s_2).$$

The probability of state $(0, Q)$ conditional on player 1 being type 0 is given by

$$\frac{p(0, Q)}{\sum_{t_2 \in T_2} p(0, t_2)} = \alpha_2(Q),$$

while the conditional probability of state $(0, 0)$ is

$$\frac{p(0, 0)}{\sum_{t_2 \in T_2} p(0, t_2)} = \lambda_2.$$

Hence, the expected payoff of type 0 of player 1 when he plays s_1 and when type 0 of player 2 plays the induced mixed strategy σ_2 is given by

$$\lambda_2 \sum_{s_2 \in S_2} \sigma_2(\delta)(s_2) u_1(s_1, s_2) + \sum_{Q \in \mathcal{Q}} \alpha_2(Q) \min_{s_2 \in Q} u_1(s_1, s_2) \quad (15)$$

We now consider two cases, first where $X_2 \setminus \tilde{X}_2 = \emptyset$ and second where $X_2 \setminus \tilde{X}_2 \neq \emptyset$. When $X_2 \setminus \tilde{X}_2 = \emptyset$, $\sigma_2(s_2) = \mu_2(\{s_2\})/\lambda_2 = \alpha_2(\{s_2\})/\lambda_2$. Hence, in this case, (15) can be written as

$$\begin{aligned} & \sum_{s_2 \in S_2} \alpha_2(\{s_2\}) u_1(s_1, s_2) + \sum_{Q \in \mathcal{Q}} \alpha_2(Q) \min_{s_2 \in Q} u_1(s_1, s_2) \\ &= \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s_1, s_2). \end{aligned} \quad (16)$$

Recall, $\lambda_2 > 0$, since we assumed $\tilde{X}_i \neq \emptyset$. When $X_2 \setminus \tilde{X}_2 \neq \emptyset$,

$$\sigma_2(s_2) = \begin{cases} (1 - \delta) \mu_2(\{s_2\})/\lambda_2 = (1 - \delta) \alpha_2(\{s_2\})/\lambda_2 & \text{if } s_2 \in \tilde{X}_2, \\ \frac{\overline{[X_2 \setminus \tilde{X}_2]}_\delta}{\lambda_2} & \text{if } s_2 \in X_2 \setminus \tilde{X}_2. \end{cases}$$

Hence, in this case, (15) can be written as

$$\begin{aligned}
& (1 - \delta) \sum_{s_2 \in \tilde{X}_2} \alpha_2(\{s_2\}) u_1(s_1, s_2) + \delta \sum_{s_2 \in X_2 \setminus \tilde{X}_2} \frac{u_1(s_1, s_2)}{|X_2 \setminus \tilde{X}_2|} + \sum_{Q \in \mathcal{Q}} \alpha_2(Q) \min_{s_2 \in Q} u_1(s_1, s_2) \\
&= \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s_1, s_2) + \delta \left[\sum_{s_2 \in X_2 \setminus \tilde{X}_2} \frac{u_1(s_1, s_2)}{|X_2 \setminus \tilde{X}_2|} - \sum_{s_2 \in \tilde{X}_2} \alpha_2(\{s_2\}) u_1(s_1, s_2) \right]
\end{aligned}$$

Meanwhile, from the hypothesis that (μ_1, μ_2, X_1, X_2) is an equilibrium of G under Knightian uncertainty, we know that every action in X_1 maximizes expected payoff for player 1. That is, any $s_1 \in X_1$ maximizes

$$\int u_1(s_1, \cdot) d\mu_2. \quad (17)$$

Recall, Lemma 1 shows,

$$\int u_1(s_1, \cdot) d\mu_2 = \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s_1, s_2).$$

Notice, since $u_1(s_1, s_2)$ is finite, corresponding to an $\epsilon_n > 0$ one may always find a $\delta_n > 0$ such that

$$\delta_n \left[\sum_{s_2 \in X_2 \setminus \tilde{X}_2} \frac{u_1(s_1, s_2)}{|X_2 \setminus \tilde{X}_2|} - \sum_{s_2 \in \tilde{X}_2} \alpha_2(\{s_2\}) u_1(s_1, s_2) \right] \leq \epsilon_n$$

Furthermore, $\lim \delta_n \rightarrow 0$ as $\lim \epsilon_n \rightarrow 0$. Hence, in the induced Bayesian game, there exists a δ_n such that given that player 2 is playing the induced mixed strategy $\sigma_2(\delta_n)$ at state $(0, 0)$, any $s_1 \in X_1$ obtains for player 1 an expected payoff at most ϵ_n away from a strategy that yields the maximum payoff. Since, by construction, $\sigma_1(\delta_n)$ puts a positive probability on s_1 iff $s_1 \in X_1$, the expected payoff to a player 1 of type 0 from playing s_1 in the induced Bayesian game, is consequently, at most ϵ_n away from a strategy that yields the maximum payoff. This completes the proof of the “only if” part of the assertion in the theorem.

To obtain proof of the “if” part of the assertion fix (μ_1, μ_2, X_1, X_2) , with $X_i \neq \tilde{X}_i$ and suppose there exists a sequence $\{(\epsilon_n, \delta_n)\}$, $\epsilon_n > 0$, $\delta_n > 0$, $\lim_{n \rightarrow \infty} (\epsilon_n, \delta_n) = (0, 0)$ such that the induced mixed strategies $(\sigma_1(\delta_n), \sigma_2(\delta_n))$ are played at state $(0, 0)$ in a strong ϵ -equilibrium of the induced Bayesian game, where $\epsilon = \epsilon_n$. Hence, for any $s_1 \in X_1, s'_1 \in S_1$, it must be true for every member of the sequence $\{(\epsilon_n, \delta_n)\}$ that,

$$\begin{aligned}
& \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s_1, s_2) + \delta_n \left[\sum_{s_2 \in X_2 \setminus \tilde{X}_2} \frac{u_1(s_1, s_2)}{|X_2 \setminus \tilde{X}_2|} - \sum_{s_2 \in \tilde{X}_2} \alpha_2(\{s_2\}) u_1(s_1, s_2) \right] \\
& \geq \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s'_1, s_2) \\
& + \delta_n \left[\sum_{s_2 \in X_2 \setminus \tilde{X}_2} \frac{u_1(s'_1, s_2)}{|X_2 \setminus \tilde{X}_2|} - \sum_{s_2 \in \tilde{X}_2} \alpha_2(\{s_2\}) u_1(s'_1, s_2) \right] - \epsilon_n \quad (18)
\end{aligned}$$

Given that $\lim_{n \rightarrow \infty} (\epsilon_n, \delta_n) = (0, 0)$, (18) implies

$$\sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s_1, s_2) \geq \sum_{A \subseteq S_2} \alpha_2(A) \min_{s_2 \in A} u_1(s'_1, s_2)$$

where $s_1 \in X_1, s'_1 \in S_1$.

Therefore, by Lemma 1

$$\int u_1(s_1, \cdot) d\mu_2 \geq \int u_1(s'_1, \cdot) d\mu_2 \text{ where } s_1 \in X_1, s'_1 \in S_1.$$

Thus every pure strategy in X_1 is a best response to the belief μ_2 in the game G . This completes the proof of the “if” part and, in turn, the proof of the theorem. ■

The second representation result shows that the Dow-Werlang equilibrium, too, may be given a representation in a Bayesian incomplete information set up. Many details of the Bayesian environment are similar to the one underlying the previous result. In particular, the type space, i.e., the nature of incompleteness of information about an opponent’s payoffs, bears similar interpretation. Here again, a player i ’s belief on any subset of the strategies played in equilibrium is “parsed” into a weight attributed to a “rational” opponent and a complementary weight distributed among “irrational” incarnations of the opponent, t_j , playing from particular subsets of strategies and choosing a strategy which minimizes player i ’s payoff. Here though, before a crazy type chooses his strategy he gets to know the outcome of any randomization resorted to by the opponent, unlike in the previous construction of the induced Bayesian game. There are important differences with respect to the nature of the representation, too. First, while in the previous case beliefs, as described in the Knightian equilibrium, could be exactly “replicated” in the Bayesian environment, in this case the replication is not exact but only approximate, though as the result shows, the approximation is arbitrarily close. Second, in the Klibanoff version, i.e., definition 1, equilibrium strategies were explicit in the description of equilibrium and the same strategies were played in the state $(0,0)$ of the induced Bayesian game. In contrast, in the Dow and Werlang version, only supports of the beliefs (and equilibrium beliefs) are made explicit in the description of equilibrium. Correspondingly, as obtained in the above result, the supports of induced mixed strategies (each an additive probability distribution) played in state $(0,0)$ of the Bayesian game are precisely the supports of the (non-additive) beliefs made explicit in the description of the Knightian equilibrium. The induced mixed strategy puts the same (to an arbitrarily close approximation) *relative* probability weight on pure actions as the weights of the Knightian equilibrium belief function on those singletons. Pure strategies which are given zero weight as singletons by the belief function but yet are in the support of the belief function receive arbitrarily small (but positive) probability weight in the strategy played at state $(0,0)$ of the induced Bayesian game.⁹ Thirdly, the equilibrium used for the representation is only a

⁹Notice though, a set of strategies, all of which receive zero weight *as singletons* in the belief function, described in the Knightian equilibrium, may well have a “large” weight in that belief function collectively (i.e., as a set). But, because the elements have zero weight as singletons, the (limiting) induced mixed strategy, in the representational equilibria of the induced Bayesian game, will assign to the set an arbitrarily small probability (the probability of the set is simply the sum of the probabilities of the elements). Of course, in the induced Bayesian game, the posterior belief (of type 0) on this set of strategies will coincide with the weight assigned by the belief function (in the Knightian equilibrium). It is just that, only a small fraction of this posterior weight is generated by the belief that the opponent will choose from this set conditional on him being the “rational” type.

(strong) ϵ -equilibrium, unlike in the earlier version where the representational equilibrium was an exact one. Of course, the ϵ may be as small as one may want, but nevertheless it is an approximate equilibrium.

Why is it that we have to work with an ϵ -equilibrium and not an exact equilibrium, as we did previously? The support X_j , described in the original equilibrium, may include strategies that receive zero weights, as singletons, in μ_j . Posterior belief of i at state $(0,0)$ in the induced Bayesian game is only an approximation to μ_j ; they differ, essentially, to the extent of δ_n , the weight the posterior puts on strategies in X_j that receive zero weights, as singletons, in μ_j . If (μ_1, μ_2, X_1, X_2) is a Knightian equilibrium of G then every member of X_i is a best response to μ_j , but not necessarily to a belief, like the posterior belief of i at state $(0,0)$, which is only an approximation to μ_j ; members of X_i are, in general, at best an approximate best response to this approximate belief.

The induced Bayesian game and the second representation result applies not only when definition 2 embodies the Dow-Werlang definition of support but also when it employs any of the other two. Recall, if the support is defined as the set consisting of all elements $s_i \in S_i$ such that $\mu_i(s_i) > 0$, i.e., $X_i \equiv \bar{X}_i$ then the definition of the induced mixed strategy is such that $\sigma_i(\delta) = \sigma_i(0)$. In this case, the equilibrium of the induced Bayesian game relevant to the representation result will be an exact equilibrium and not a (non-trivial) ϵ -equilibrium. It is important to note that the different notions of support will imply different behavior at the state $(0,0)$ in the relevant equilibrium of the induced Bayesian game. The support of the induced mixed strategy in the equilibrium of the induced Bayesian game coincides *exactly* with what is defined to be the support in the description of the original Knightian equilibrium; if the definition of support in the Knightian equilibrium changes so does the support of the relevant induced mixed strategy.

4.3 Common Belief of Game G

Take an equilibrium under Knightian uncertainty of the game G per definition 2 wherein the equilibrium belief functions are strictly non-additive. Consider, the corresponding induced Bayesian game and an ϵ -equilibrium in induced mixed strategies of this game. Suppose both players are of type 0. Then, in no ϵ -equilibrium under consideration will it be common knowledge between the players that the underlying game being played is the benchmark game G . Indeed, the greater the uncertainty aversion embodied in the belief functions, the greater will be the deviation from common knowledge. Only when the beliefs functions are conventional, *additive* probability measures will there be common knowledge of the game being played. Proposition 2, below, states this claim more precisely.

Proposition 2 *Suppose (μ_1, μ_2, X_1, X_2) is an equilibrium of G under Knightian uncertainty according to definition 2. Consider the induced Bayesian game and the sequence of strong ϵ -equilibria of the induced Bayesian game corresponding to a sequence $\{(\epsilon_n, \delta_n)\}$, $\lim_{n \rightarrow \infty} (\epsilon_n, \delta_n) = (0, 0)$, such that the induced mixed strategies $(\sigma_1(\delta_n), \sigma_2(\delta_n))$ are played at state $(0, 0)$ and $\epsilon = \epsilon_n$. Then, at all such ϵ -equilibria of the induced Bayesian game, at*

state $(0,0)$ there is common λ -belief that the game G is being played, where

$$\lambda = \min \left\{ \sum_{s_1 \in S_1} \mu_1(\{s_1\}), \sum_{s_2 \in S_2} \mu_2(\{s_2\}) \right\}.$$

For any $\lambda' > \lambda$, it is not common λ' -belief that the game G is being played.

If at least one of the belief functions (μ_1, μ_2) is non-additive, we have $\lambda < 1$, so that there is a failure of common 1-belief of the game being played. In general, the lower is the weight given to the singleton sets $\{s_i\}$ by these belief functions, the greater will be the departure from common belief. The proof of this result is very similar to that for Proposition 1 and, hence, is omitted.

5 Conclusion

The exercise in this paper shows that it is possible to represent an equilibrium under Knightian uncertainty (with or without preference for randomization) as an equilibrium of an incomplete information game wherein players have incomplete information about their opponent's payoff functions and information structure of the game.¹⁰ The incomplete information game constructed for the representation has both standard and non-standard features. In particular, it is standard with respect to the feature that players have Bayesian beliefs. On the other hand, the way we model how players think about contingent states is not standard. In the incomplete information game constructed, at an equilibrium relevant to the representation, the game G is not common knowledge: player i believes there can be contingencies where the game is not as given in G . The non-standard features involve i 's beliefs about the contingent games. An agent i believes that in these games, one, he is facing an opponent who is not committed to a particular strategy but is choosing from a restricted set of strategies; two, the opponent will choose a strategy that minimizes i 's payoffs; three, the opponent is informed, to a varying degree depending upon the notion of Knightian equilibrium invoked in the original game, about the strategy chosen by i . Indeed, these three non-standard features of the representative induced Bayesian game “explain” why we may observe behavior in a Knightian equilibrium that cannot be observed in a standard complete or incomplete information game. The distinctive consequent behavior is different depending upon the nature of the third non-standard feature incorporated, more specifically, the information the crazy type is assumed to have about its opponent's strategy when

¹⁰Footnote 4 mentioned the possibility of using the Ghirardato-Le Breton decomposition to extend the analysis to general capacities. A referee suggests a construction of an induced Bayesian game along the following lines. The construction is not only ingenious, but also most instructive about the general nature of complications involved. First recall, the Ghirardato-Le Breton decomposition involves summing over collections of non-nested sets from each of which a set to minimize over is selected by maximizing (as opposed to simply summing over sets, each of which is used to minimize over, as in Lemma 1). When we apply this to an equilibrium in general capacities, in the “artificial” states of the world, the extensive-form would involve an additional (third) player who shares the first-mover's payoffs and also observes the first-mover's strategy. The first-mover would pick a strategy, then this additional player would be allowed to choose (from limited options) the set of strategies available to the “malicious” player who (as in the present construction) has opposite payoffs from the first-mover, finally the malicious player would move. Thus a “type” would determine the set of action sets the additional player could select from (with belief functions this set of sets is always a singleton, so the additional player does not appear in the present paper).

it makes its choice. That in turn depends on the notion of Knightian equilibrium invoked in the original game: corresponding to the Klibanoff version, the crazy type only knows the randomized strategy but not the outcome of the randomization and the consequence of the non-standard features is that there is a strict preference for randomization in equilibrium; corresponding to the Dow-Werlang version, the crazy type is assumed to know the outcome of the randomization, and here the non-standard features may lead to an equilibrium strategy that is dominated by a mixed strategy in a standard game. (The reader will recall these consequences were illustrated in examples I and II, respectively.)

Since beliefs and behavior are both explicitly described in the statement of the Klibanoff version of Knightian equilibrium, the Bayesian equilibrium representing this version replicates both belief and behavior described in the original equilibrium. Beliefs are replicated in the sense that the equilibrium posterior beliefs on strategies in the Bayesian game are the basic probability assignments constituting the belief function described in the Knightian equilibrium. The behavior is replicated in the sense that the strategy that is actually played in the Bayesian game, i.e., the strategy at the only state which actually occurs, state (0,0), is the strategy described in the Knightian equilibrium. The Dow-Werlang version of Knightian equilibrium is described only in terms of beliefs and supports of these beliefs. The Bayesian equilibrium representing the Dow-Werlang equilibrium replicates the beliefs (though only up to an arbitrarily close approximation) in much the same way as with the Klibanoff version. The representation also replicates the support, in the sense that the support of the induced mixed strategy $\sigma_i(\delta_n)$, the strategy played at state (0,0) of the induced Bayesian game, is the same as the support X_i described in the original equilibrium.

Acknowledgements

The authors are grateful to the editor, Chris Shannon, and three referees for their many constructive suggestions and insightful comments. Their feedback has improved the content and exposition very significantly. We also thank J.Eichberger, L.Epstein, P. Ghirardato, D.Kelsey, P. Klibanoff, M. Marinacci, S. Morris and P.Wakker for their comments on earlier drafts.

Sujoy Mukerji may be contacted at sujoy.mukerji@economics.ox.ac.uk and Hyun Song Shin may be contacted at h.s.shin@lse.ac.uk.

References

- [1] Choquet, G. (1954): "Theory of Capacities," *Ann. Inst. Fourier* (Grenoble), 5, 131-295.
- [2] Dempster, A. P. (1967): "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals of Mathematical Statistics*, 38, 325-339.
- [3] Dempster, A. P. (1968): "A Generalization of Bayesian Inference," *Journal of the Royal Statistical Society, Series B*, 205-247.
- [4] Dow J. and S. Werlang (1994): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64, 305-324.

- [5] Dubois, D. and H. Prade (1985): "Evidence Measures Based on Fuzzy Information," *Automatica*, 21, 547-562.
- [6] Eichberger, J. and D. Kelsey (1996) "Uncertainty Aversion and Preference for Randomization," *Journal of Economic Theory*, 71, 31-43.
- [7] Eichberger, J. and D. Kelsey (2000) "Non-Additive Beliefs and Strategic Equilibria," *Games and Economic Behaviour*, 30, 183-215
- [8] Epstein, L. G. (1997): "Preference, Rationalizability and Equilibrium," *Journal of Economic Theory*, 73, 1-29.
- [9] Epstein, L. G. and T. Wang (1996) "Beliefs about Beliefs without Probabilities" *Econometrica*, 64, 1343-73.
- [10] Ghirardato, P. (2001): "Coping with Ignorance: Unforeseen Contingencies and Non-additive Uncertainty," *Economic Theory*, 17:247-276.
- [11] Ghirardato, P. and M. LeBreton (1998): "Choquet Rationality," *mimeo. Division of Humanities and Social Sciences, California Institute of Technology*.
- [12] Ghirardato, P. and M. LeBreton (2000): "Choquet Rationality," *Journal of Economic Theory*, 90, 277-85.
- [13] Gilboa, I. and D. Schmeidler (1994): "Additive Representations of Non-additive Measures and the Choquet Integral," *Annals of Operations Research*, 51, 43-65.
- [14] Gilboa, I. and D. Schmeidler (1989): "Maxmin expected utility with a non-unique prior," *Journal of mathematical Economics*, 18, 141-153.
- [15] Groes, E., H. J. Jacobsen, B. Sloth, and T. Tranaes (1998): "Nash Equilibrium in Lower Probabilities," *Theory and Decision*, 44, 37-66.
- [16] Haller, H. (1995): "Non-additive Beliefs in Solvable Games," *mimeo., Virginia Polytechnic Institute*, Blacksburg, VA.
- [17] Jaffray, J.-Y. (1989): "Linear Utility Theory for Belief Functions," *Operations Research Letters*, 8, 107-112.
- [18] Jaffray, J.-Y., P. Wakker (1994): "Decision Making with Belief Functions: Compatibility and Incompatibility with the Sure-thing principle," *Journal of Risk and Uncertainty*, 8, 255-271.
- [19] Klibanoff, P., (1996): "Uncertainty, Decision and Normal Form Games," *mimeo. Northwestern University*. (Revised version of an earlier mimeo. dated (1993))
- [20] Klibanoff, P., (2001): "Stochastically independent randomization and uncertainty aversion," *Economic Theory*, 18, 605-620.
- [21] Knight, F. H (1921): *Risk, Uncertainty and Profit*, Houghton Mifflin, Boston.

- [22] Lipman, B. (1999): "Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality." *Review of Economic Studies*, 66, 339-61.
- [23] Lo, K. C. (1996): "Equilibrium in Beliefs Under Uncertainty," *Journal of Economic Theory*, 71, 443-84.
- [24] Marinacci, M. (2000): "Ambiguous Games," *Games and Economic Behavior* 31, 191-219.
- [25] Mertens, J.-F. and S. Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information" *International Journal of Game Theory*, 14, 1-29.
- [26] Monderer D. and Samet P. (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170-190.
- [27] Mukerji, S. (1995): "A theory of play for games in strategic form when rationality is not common knowledge," *University of Southampton discussion paper* 9519.
- [28] Mukerji, S. (1997): "Understanding the Nonadditive Probability Decision Model," *Economic Theory*, 9, 23-46.
- [29] Radner, R. (1980): "Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22, 136-54.
- [30] Ryan, M. J. (1999): "What do Uncertainty-Averse Decision-Makers Believe?" *mimeo. University of Auckland*.
- [31] Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571-587.
- [32] Shafer, G. (1976): *A Mathematical Theory of Evidence*, Princeton University Press, New Jersey.
- [33] Shafer, G. (1981): "Constructive Probability" *Synthese*, 48, 1-59.
- [34] Smets, P. (1981): "The Degree of Belief in a Fuzzy Event," *Information Science*, 25.
- [35] Wasserman, L. A. (1990): "Belief Functions and Statistical Inference," *The Canadian Journal of Statistics*, 18, 183-196.