



Ambiguous act equilibria

Sophie Bade^{*,1}

Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, D-53113 Bonn, Germany

ARTICLE INFO

Article history:

Received 16 July 2008

Available online 10 May 2010

JEL classification:

C72

D81

Keywords:

Uncertainty aversion

Nash equilibrium

Ambiguity

ABSTRACT

A game-theoretic framework that allows for explicitly randomized strategies is used to study the effect of ambiguity aversion on equilibrium outcomes. The notions of “independent strategies” as well as of “common priors” are amended to render them applicable to games in which players lack probabilistic sophistication. Within this framework the equilibrium predictions of two-player games with ambiguity-averse and with ambiguity-neutral players are observationally equivalent. This equivalence result does not extend to the case of games with more than two players. A translation of the concept of equilibrium in beliefs to the context of ambiguity aversion yields substantially different predictions – even for the case with just two players.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

The Ellsberg paradox shows that people treat risky situations in which they know the odds of all relevant outcomes differently from ambiguous situations in which they can only guess these odds. It has inspired a large range of different generalizations of expected utility theory that can accommodate for the apparent aversion to ambiguity. This branch of decision theory continues to thrive.² If anything we might expect that ambiguity aversion is more relevant for strategic decision making than it is for single person decision making as people are quite generally harder to predict than natural processes.

It is surprising, then, that the literature on games among players that are ambiguity-averse stayed comparatively small. A review of some applications can be found in Mukerji and Tallon (2004). It is important to mention that in many of the applications of such games the players are assumed to be uncertain about the environment rather than about each other's strategies; see Bade (2007), Levin and Ozdenoren (2004), Bose et al. (2006). In contrast, the present paper concerns the case in which players are assumed to be uncertain about each other's strategies.

To do so, the present paper takes up the interpretation of mixed-strategy equilibrium according to which it assumed that players explicitly randomize over actions. It replaces the assumption that players are expected utility maximizers with the assumption that players are averse to ambiguity. It is shown that for the class of two-player games the equilibrium predictions are not affected by such a change in the assumption on the players' attitudes towards subjective uncertainty.

To accommodate different attitudes towards ambiguity I introduce the assumption that players can choose to play “ambiguous strategies”. The players in the present approach cannot only base their strategies on objective randomization devices

* Fax: +49 228 9141670.

¹ Thanks to Kalyan Chatterjee, Christoph Engel, Ed Green, Martin Hellwig, Andrei Karavaev, Vijay Krishna, Matthias Lang, Kin Chung Lo, Yusufcan Masatlioglu, Efe Ok, Emre Ozdenoren, Joris Pinkse, Ludovic Renou, Muhamet Yildiz and the seminar participants at Bocconi, Bielefeld, DUKE/UNC, Collegio Carlo Alberto, the Game Theory Festival at Stony Brook, Leicester, LUISS, Oxford, Paris and the RUD conference in Tel Aviv.

² For a review on the experimental evidence, see Camerer and Weber (1992). Some of the seminal contributions to the decision theory of ambiguity aversion are Schmeidler (1989), Gilboa and Schmeidler (1989), and Bewley (2002); for a very recent treatment that generalizes many of the available representations, see Cerreia-Vioglio et al. (2008).

such as dices or roulette wheels, they can also use subjective randomization devices. Players can, for instance, base the choice of their strategy on the draws from an Ellsberg urn, on a horse race or even on their feelings. This generalization of the players' strategy spaces makes it possible to endow players with different attitudes towards ambiguity.

The generalization suggested so far actually overshoots the mark of parsimonious deviation that allows us to identify the effect of the players' ambiguity attitude on equilibrium predictions: the introduction of *all* kinds of subjective randomization devices not only allows for different attitudes towards ambiguity, but also for correlation devices and wildly diverging beliefs. Sections 3.1 and 3.2 are devoted to tackling these two issues. This is somewhat harder than one might initially think, as I cannot rely on probabilistic beliefs on the state space to define the notions of "independent strategies" and "common priors". Once these hurdles are taken, I use the standard Nash equilibrium concept to define an Ambiguous Act Equilibrium in Section 4 as a profile of ambiguous act strategies such that no player has an incentive to deviate given all other players' strategies.

The definition of an Ambiguous Act Equilibrium can be seen as an answer to the first of three questions that Mukerji and Tallon (2004) identified as the guiding questions in research on game theory with ambiguity-averse players. In their review of applications of David Schmeidler's concept of uncertainty aversion, they describe these three questions as follows: "(1) ... how should solution concepts ... be defined?; (2) questions about the general behavioral implications of the new solution concepts; (3) questions about insights such innovations might bring to applied contexts". My main contribution to the second question is a negative one: in the main result of the present paper (Theorem 1), I show that the ambiguous act equilibria and the mixed strategy equilibria of a two-player game are observationally equivalent. In Sections 6 to 10, I ask and answer the following questions: are there any action profiles that can arise in mixed-strategy equilibrium, but would never arise in ambiguous act equilibrium? Conversely, are there any action profiles that are consistent with the assumption of equilibrium play by ambiguity-averse players, but are inconsistent with mixed-strategy equilibrium? Can the observation of an action profile tell us whether the players are ambiguity-neutral or averse? For games between two players with transitive and monotone preferences, the answer is negative. In such games, the set of all ambiguous act equilibria is observationally equivalent to the set of all mixed strategy equilibria, in the sense that any action profile that is consistent with equilibrium play among uncertainty-averse players is consistent with equilibrium play among uncertainty-neutral players. In the present context, the answer to Mukerji and Tallon's questions (numbers (2) and (3)) is that the general behavioral implications of uncertainty aversion are not different from the general behavioral implications of expected utility maximization and consequently, that there is little hope for new insights in applied contexts.

I go on to show in Sections 11.2 and 11.3 that the prospects for different behavioral implications of different ambiguity attitudes do not look as bleak once we leave the setup of this paper. I show, in particular, that observational difference between the equilibrium play of ambiguity-averse players can either arise when there are more than two players or when the equilibrium concept proposed here is replaced by an equilibrium concept that extends the notion of equilibrium in beliefs to the context of ambiguity-averse players. I also show that some of the existing results of behavioral difference in the literature on games with ambiguity-averse players mainly derive from the assumption that allows for players' beliefs to diverge "extremely". I juxtapose my definition of ambiguous act equilibrium with existing equilibrium concepts in Section 5, and I compare my results to some of the equilibrium predictions in the literature on games with ambiguity-averse players in Section 11.4.

2. General ambiguous games

Two-player normal form games (in mixed strategies) are denoted by $G = (A, \succsim)$ with $A := (A_1, A_2)$, the players' sets of pure strategies, which are assumed to be finite. The players preferences are summarized by $\succsim := (\succsim_1, \succsim_2)$. They are defined on the set $\Delta(A)$, where $\Delta(N)$ denotes the set of all lotteries on any (finite) N . A *general ambiguous act extension* of G is a game $G' = (\Omega, A, \succsim')$ with the following interpretations of Ω and \succsim' : the state space Ω can be represented as the product $\Omega_1 \times \Omega_2$ of two finite "state spaces" Ω_1 and Ω_2 . The preferences $\succsim' = (\succsim'_1, \succsim'_2)$ are defined over all acts $f : \Omega \rightarrow \Delta(A)$; the restriction of these preferences to the set of all constant acts equals the preferences in the original game G , formally $\succsim'_i |_{\Delta(A)} = \succsim_i$ for $i = 1, 2$. The set of player i 's strategies in the ambiguous act extension G' is the set of all acts $f_i : \Omega_i \rightarrow \Delta(A_i)$. A strategy profile (f_1, f_2) induces an act $f : \Omega \rightarrow \Delta(A)$ with $f(\omega)(a) = f_1(\omega_1)(a_1)f_2(\omega_2)(a_2)$ for all $a \in A$. So the probability that an action-profile $a = (a_1, a_2)$ is being played in state ω is determined as the product of the probabilities that the two players play action a_i in state ω . I denote the act induced by a strategy profile (f_1, f_2) as well as the strategy profile itself by f . A game G' is called a *general ambiguous game*, if it is a general ambiguous act extension of some game G .

The assumption that player i 's action space consists of all acts $f_i : \Omega_i \rightarrow \Delta(A_i)$ implies an assumption on player i 's knowledge: player i can only base his actions on the i 'th component of the state since he only knows this aspect of the state. Player i 's knowledge can be described by the event algebra \mathcal{S}_i on Ω such that $E_i \in \mathcal{S}_i$ if $\omega \in E_i$ implies that $(\omega_i, \omega'_{-i}) \in E_i$ for all $\omega'_{-i} \in \Omega_{-i}$.³ This assumption is without loss of generality as I could have defined player i 's strategy space equivalently as the set of all \mathcal{S}_i -measurable acts $f : \Omega \rightarrow \Delta(A_i)$. It is convenient to define the events $\omega_i^* := \{\omega \mid \omega_i = \omega_i^*\}$.

³ I follow the usual convention and define $x_{-i} = x_2$ if $i = 1$ and $x_{-i} = x_1$ otherwise.

The assumption on strategy spaces also implies that players have access to objective randomizing devices that can generate any lottery on the action spaces A_i . A player that can choose any $f_i : \Omega_i \rightarrow \Delta(A_i)$ is free to generate his strategic choices using roulette wheels, dice, objective computer generators or similar things. In the games under study, players are equally free to base their choices on their mood of the day, or on any other subjective random device to which they have access. The definition of a game does not depend on this notion of an act.

Let me note as an aside that games could also be defined for an environment with Savage acts $f : \Omega \rightarrow A$, in which no objective lotteries are assumed. This could be interesting as the empirical evidence on mixed strategies suggests that “normal” people are not able to play mixed strategies. Chiappori et al. (2002), Palacios-Huerta (2003), and Walker and Wooders (2001) therefore use “abnormal” people, namely athletes, to test the predictions of mixed-strategy equilibrium. Games played by “normal” people could be studied in a framework in which strategies are acts $f : \Omega_i \rightarrow A_i$.

2.1. Notation

I use the letters f, g, f_i, g_i to denote acts. Lotteries on action profiles and action spaces are denoted by $p, q \in \Delta(A)$ or $p_i, q_i \in \Delta(A_i)$, respectively. As a shorthand, I denote a constant act f with $f(\omega) = p$ for all $\omega \in \Omega$ and some $p \in \mathcal{P}(A)$ directly by p (and accordingly for f_i). Degenerate lotteries, that is, lotteries $p \in \Delta(A)$ and $p_i \in \Delta(A_i)$ such that $p(a) = 1$ for some a or $p_i(a_i) = 1$ for some a_i , are denoted a or a_i . Finally, constant acts with $f(\omega) = a$ or $f_i(\omega_i) = a_i$ for all $\omega \in \Omega$ or $\omega_i \in \Omega_i$ are denoted by a and a_i , respectively. Constant acts a correspond to pure strategy profiles, constant acts a_i correspond to pure strategies. Constant acts p and p_i correspond to mixed-strategy profiles and mixed strategies, respectively. In short, pure and mixed strategies are naturally embedded in the framework of general ambiguous games.

The mixture $\alpha f + (1 - \alpha)g$ of two acts f, g is defined componentwise, meaning that $(\alpha f + (1 - \alpha)g)(\omega)(a) = \alpha f(\omega)(a) + (1 - \alpha)g(\omega)(a)$ for all $a \in A$ and all $\omega \in \Omega$. The complement of an event $E \subset \Omega$ is denoted by \bar{E} . For any three acts f, g, h and any mutually exclusive events $E, F \subset \Omega$, I define the act $(f : E, g : F, h : \overline{E \cup F})$ by

$$(f : E, g : F, h : \overline{E \cup F})(\omega) = \begin{cases} f(\omega) & \text{for } \omega \in E, \\ g(\omega) & \text{for } \omega \in F, \\ h(\omega) & \text{otherwise.} \end{cases}$$

An event E is *i-null* if $(f : \bar{E}, g : E) \sim_i (f : \bar{E}, h : E)$ holds for all acts $f, g, h : \Omega \rightarrow \Delta(A)$. If an event is not *i-null*, then it is called *i-non-null*.

3. Parsimony

The goal of the present study is to see how the equilibrium predictions for a game change when the assumption of players maximizing expected utility is replaced by the assumption of ambiguity-averse players. To reach this goal, the games studied here should differ from standard games in mixed strategies in only this one respect. In other words, the games studied here should reduce to standard games in mixed strategies when all players’ preferences have expected utility representations $U = (u, \pi)$ consisting of utility functionals $u : \Delta(A) \rightarrow \mathbb{R}$ and priors $\pi \in \Delta(\Omega)$ such that $f \succsim g$ holds for two acts $f, g : \Omega \rightarrow \Delta(A)$, if and only if $U(f) = \sum_{\omega \in \Omega} \pi(\omega)u(f(\omega)) \geq U(g) = \sum_{\omega \in \Omega} \pi(\omega)u(g(\omega))$. The general ambiguous act extensions defined above are too general for the purpose of this study. They not only allow for various ambiguity attitudes; they also allow for correlation devices and wildly diverging beliefs.

In fact the notion of a general ambiguous game corresponds to the definition of a game that Aumann (1974) uses in his introduction of the concept of correlated equilibrium. Aumann starts out with the same general definition of a game and goes on to impose an expected utility representation. The present project can be seen as complementary to Aumann’s: how would the set of equilibria change if we dropped the assumption that players are expected utility maximizers, but retained the assumption that players cannot rely on any correlation devices? With the goal of the most parsimonious deviation from standard theory that allows for the introduction of a new aspect, I should proceed by imposing that the strategies of all players are independent and that all players share some common belief on the underlying state space. In the following two subsections, I develop notions of strategic independence and common beliefs that can be applied to games that do not allow for an expected utility representation of preferences. This allows me to define the notion of ambiguous act extensions, as general ambiguous act extensions that satisfy the requirements of strategic independence and common beliefs.

3.1. Independent strategies

Strategies are considered independent, if the randomization devices (namely state spaces Ω_1 and Ω_2) on which players can condition their strategies are stochastically independent. The definition of stochastic independence advocated here builds on the idea that the value of a mixed strategy p_i should not depend on the event $E_i \subset \Omega_i$ in which it is played. It is helpful to consider the standard case in which preferences have an expected utility representation, as defined above, to develop and understand this definition. Consider an act $f = (f_1, f_2)$ and assume that $\pi(\omega_1) > 0$ for all $\omega_1 \in \Omega_1$. Observe that $U(f)$ can alternatively be expressed as the weighted sum of conditional expected utilities:

$$U(f) = \sum_{\omega_1 \in \Omega_1} \pi(\omega_1) \sum_{\omega_2 \in \Omega_2} u(f_1(\omega_1) \times f_2(\omega_2)) \frac{\pi(\omega_1 \cap \omega_2)}{\pi(\omega_1)}.$$

Observe that Ω_1 and Ω_2 are stochastically independent according to the standard definition, if and only if $\frac{\pi(\omega_1 \cap \omega_2)}{\pi(\omega_1)} = \pi(\omega_2)$ for all $\omega \in \Omega$. This, in turn, implies that the conditional value of a mixed strategy p_1 does not depend on the event in which it is played; in other words, the value of $U((p_1 \times f_2) | \omega_1) = \sum_{\omega_2 \in \Omega_2} u(p_1 \times f_2(\omega_2)) \frac{\pi(\omega_1 \cap \omega_2)}{\pi(\omega_1)}$ does not depend on ω_1 as $\frac{\pi(\omega_1 \cap \omega_2)}{\pi(\omega_1)}$ is equal to $\pi(\omega_2)$ no matter which ω_1 we choose.

So in a sense, with independent state spaces Ω_1 and Ω_2 , strategies become “exchangeable”: the mixture p_1 yields the same conditional value whether it is played in the event ω_1 or in the event ω'_1 . This, however, does not mean that for independent strategies players are indifferent between any acts $((p_1 : E, q_1 : \bar{E}), f_2)$ and $((q_1 : E, p_1 : \bar{E}), f_2)$. If the two events E, \bar{E} do not have the same size, the acts might yield a different expected utility. The definition of stochastic independence proposed here builds on the idea of event exchangeability corrected for the different size or importance of different events.

Definition 1. Take a general ambiguous game $G' = (\Omega, A, \succsim')$. Then Ω_1 and Ω_2 are called k -independent if for all strategy profiles $f : \Omega \rightarrow \Delta(A)$, all mutually exclusive k -non-null events $E, F \subset \Omega_i$ and all mixtures $p_i, q_i \in \Delta(A_i)$ there exists an $\alpha \in [0, 1]$ such that either

$$((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i}) \sim'_k ((\alpha p_i + (1 - \alpha)q_i : E, p_i : F, f_i : \overline{E \cup F}), f_{-i})$$

or

$$((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i}) \sim'_k ((q_i : E, \alpha p_i + (1 - \alpha)q_i : F, f_i : \overline{E \cup F}), f_{-i}).$$

The parameter α acts as a corrective for the possibly different “size” of E and F . If players have a prior on Ω , these “sizes” can straightforwardly be identified with their probability, as can be seen in the proof of the following lemma. When agents do not base their decisions on a (single) prior, this identification fails and the value of the corrective parameter α is therefore allowed to adjust endogenously.

Lemma 1. If a player’s preferences have an expected utility representation, then the behavioral notion of independence given in Definition 1 coincides with the standard notion of independence.

Proof. Assume that Ω_1 and Ω_2 are independent according to the standard definition. Then we have for any act $((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i})$ with $\pi(E) \geq \pi(F)$

$$\begin{aligned} &U((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i}) \\ &= \pi(E) \sum_{\omega_2 \in \Omega_2} u(p_i \times f_2(\omega_2)) \frac{\pi(E \cap \omega_2)}{\pi(E)} + \pi(F) \sum_{\omega_2 \in \Omega_2} u(q_i \times f_2(\omega_2)) \frac{\pi(F \cap \omega_2)}{\pi(F)} + R \\ &= \pi(E) \sum_{\omega_2 \in \Omega_2} u(p_i \times f_2(\omega_2)) \pi(\omega_2) + \pi(F) \sum_{\omega_2 \in \Omega_2} u(q_i \times f_2(\omega_2)) \pi(\omega_2) + R \\ &= \pi(F) \sum_{\omega_2 \in \Omega_2} u(p_i \times f_2(\omega_2)) \pi(\omega_2) + \pi(E) \sum_{\omega_2 \in \Omega_2} u((\alpha q_i + (1 - \alpha)p_i) \times f_2(\omega_2)) \pi(\omega_2) + R \\ &= U(((p_i : F, \alpha q_i + (1 - \alpha)p_i : E, f_i : \overline{E \cup F}), f_{-i})) \end{aligned}$$

with $R = \sum_{\omega_2 \in \Omega_2, \omega_1 \notin E \cup F} u(f_1(\omega_1) \times f_2(\omega_2)) \frac{\pi(\omega_1 \cap \omega_2)}{\pi(\omega_1)}$ and $\alpha := \frac{\pi(F)}{\pi(E)}$.

Conversely, assume that Ω_1 and Ω_2 are not independent according to the standard definition. Then, there exist non-null events $E \subset \Omega_1$ and $F \subset \Omega_2$ such that $\pi(E \cap F) > \pi(E)\pi(F)$, $\pi(\bar{E} \cap \bar{F}) > \pi(\bar{E})\pi(\bar{F})$, $\pi(\bar{E} \cap F) < \pi(\bar{E})\pi(F)$ and $\pi(E \cap \bar{F}) < \pi(E)\pi(\bar{F})$. Define a game in which player 1 can choose among the strategies T and B , and player 2 can choose from L and R such that the payoffs to player one are as follows: $u(T \times L) = 1$, $u(B \times R) = \frac{\pi(F)}{\pi(E)}$ and $u(T \times R) = u(B \times L) = 0$. Define strategy $f_2 := (L : F, R : \bar{F})$. Observe that

$$U_1((T : E, B : \bar{E}) \times f_2) = \pi(E \cap F) + \pi(\bar{E} \cap \bar{F}) \frac{\pi(F)}{\pi(E)} > \pi(E)\pi(F) + \pi(\bar{E})\pi(\bar{F}) \frac{\pi(F)}{\pi(E)} = \pi(F)$$

while on the other hand

$$\begin{aligned}
U_1((B : E, \alpha T + (1 - \alpha)B : \bar{E}) \times f_2) &= \pi(E \cap \bar{F}) \frac{\pi(F)}{\pi(\bar{F})} + \alpha \pi(\bar{E} \cap F) + (1 - \alpha) \pi(\bar{E} \cap \bar{F}) \frac{\pi(F)}{\pi(\bar{F})} \\
&= \alpha \left(\pi(E \cap \bar{F}) \frac{\pi(F)}{\pi(\bar{F})} \right) + \pi(\bar{E} \cap F) + (1 - \alpha) \pi(F) \\
&\leq \alpha \left(\pi(E) \pi(\bar{F}) \frac{\pi(F)}{\pi(\bar{F})} + \pi(\bar{E}) \pi(F) \right) + (1 - \alpha) \pi(F) \\
&= \pi(F), \\
U_1((\alpha T + (1 - \alpha)B : E, T : \bar{E}) \times f_2) &= \alpha \pi(E \cap F) + (1 - \alpha) \pi(E \cap \bar{F}) \frac{\pi(F)}{\pi(\bar{F})} + \pi(\bar{E} \cap F) \\
&= \alpha (\pi(E \cap F) + \pi(\bar{E} \cap F)) + (1 - \alpha) \left(\pi(E \cap \bar{F}) \frac{\pi(F)}{\pi(\bar{F})} + \pi(\bar{E} \cap F) \right) \\
&\leq \alpha \pi(F) + (1 - \alpha) \left(\pi(E) \pi(\bar{F}) \frac{\pi(F)}{\pi(\bar{F})} + \pi(\bar{E}) \pi(F) \right) \\
&= \pi(F)
\end{aligned}$$

for all $\alpha \in [0, 1]$. So there does not exist an α to satisfy the demanded indifference. \square

The intuition behind the latter part of the proof translates to preferences that are not representable by expected utility functions. Suppose player 2 conditions his strategy on some event F . Consider player 1's best reply under the assumption that he can condition his strategy on an event E which is correlated with F . Given this correlation, player 1 will choose different strategies for the case that E does happen and the case that E does not happen. The correlation also implies that switching these two choices with each other will lower player 1's utility as he would now play the strategy which fits player 2's strategy in the event F best in the event \bar{E} which is *negatively* correlated with F and similarly for \bar{F} . A simple correction for the size of the events (the parameter α) will not suffice to elevate player 1's utility to the level achieved when playing the best response.

I am aware of three alternative behavioral definitions of independence in the literature by Brandenburger et al. (1991), Klibanoff (2001) and Bade (2008). None of the three definitions builds on the idea of event exchangeability. Brandenburger, Blume, and Dekel's definition of independence uses updated preferences to say that an event E is independent of an algebra \mathcal{S} if preferences over two \mathcal{S} -measurable acts f and g never change when learning (and updating with respect to) the independent event E . Defining stochastic independence through updating in the present context would be unpractical as there is little agreement on how to update preferences of decision-makers whose preferences are not representable by expected utility functions. Klibanoff (2001) defines two algebras \mathcal{S}_1 and \mathcal{S}_2 independent, if an agent is indifferent between any σ_i -measurable act f and any corresponding act f^E , according to which f is played only if some event E in σ_{-i} obtains and the certainty equivalent of f is paid otherwise. Bade (2008) characterizes independent events for preferences that can be represented by maxmin utilities following Gilboa and Schmeidler (1989) and following various modifications of Klibanoff's notion of independence. The question how the definition of stochastic independence given above relates to the definitions provided by Klibanoff (2001) and Bade (2008) remains open.

3.2. Agreement

In addition, for correlation devices, general ambiguous games allow for wildly diverging beliefs. To see this, take the following game between two players that can either travel to Rome or to Paris. Consider a state space $\Omega_1 = \{\omega_1, \omega'_1\}$ and a strategy profile which prescribes that player 1 goes to Rome, if and only if state ω_1 occurs. Finally, specify beliefs such that player 1 is sure that ω_1 occurs and player 2 is sure that ω'_1 occurs. We can find a pure strategy equilibrium even if we specify preferences such that one player prefers meetings to being in different cities, whereas the other has the inverse preference.

The constructed game, together with its unappealing equilibrium, fly in the face of the common prior assumption. In keeping with the goal of a parsimonious deviation from the theory of mixed-strategy equilibrium, I need to impose a condition that would eliminate such games from consideration. The difficulty lies in the fact that the common prior assumption, just like strategic independence, is defined in terms of the players' expected utility representations. When players have no priors on the state space, they cannot be common.

To derive a notion of common beliefs, observe that the set of i -null states can be interpreted as the set of states that are "never" going to happen following player i 's belief. If player i prefers an act f to an act g , even though these two acts only differ on a state ω , player i should better believe that this event could possibly happen. Conversely, if player i is indifferent between all acts that differ only on a state ω , this event is irrelevant for player i 's payoff; he might as well think that this event will never happen. Instead of "common priors", I propose to require that an event is 1-null if and only if it is 2-null. This assumption can be interpreted as the assumption that players agree on the set of states that could possibly happen. Observe that this requirement is generally strictly weaker than the requirement of "common priors": if the players'

preferences are representable by expected utilities, players agree on null events if and only if they agree on probability zero events. Given the agreement on null-events, one can sensibly define an event as *null* if and only if it is null for player 1 or 2 and as *non-null* otherwise. The set of all non-null states of player i is denoted by Ω_i^m .

4. Ambiguous games and ambiguous act equilibria

A general ambiguous game $G' = (\Omega, A, \succ')$ is considered an *ambiguous game*, if Ω_1 and Ω_2 are both 1- and 2-independent and if an event is 1-null if and only if it is 2-null. Analogously, a general ambiguous act extension $G' = (\Omega, A, \succ')$ of $G = (A, \succ)$ is considered an *ambiguous act extension* of $G = (A, \succ)$ if it is an ambiguous game. Ambiguous games and ambiguous act extensions do not reduce to standard games in mixed strategies or standard mixed extensions of normal form games, when imposing the requirement that all players' preferences are representable by expected utilities, since the requirement that players agree on non-null events is not sufficient to establish common priors. Remember, though, that the main result of this paper (Theorem 1) is a negative one: it states that the equilibrium play of ambiguity-averse and ambiguity-neutral agents is observationally equivalent. Consequently, the main result of this paper is *strengthened* by the fact that ambiguous games with expected utility representations are a superclass of games in mixed strategies: no observational difference can be derived from varying the degree of ambiguity aversion *and* the degree of disagreement between players – as long as they agree on null events. This finally sets the stage for the definition of ambiguous act equilibria as Nash equilibria of ambiguous act extensions.

Definition 2. Let $G = (A, \succ)$. A strategy profile f is called an *ambiguous act equilibrium* (AAE) of G , if G has an ambiguous act extension, $G' = (\Omega, A, \succ')$ such that

$$f \succ'_i (f'_i, f_{-i}) \quad \text{for all } f'_i: \Omega_i \rightarrow \Delta(A_i) \text{ for } i = 1, 2.$$

The notion of ambiguous act equilibrium straightforwardly extends the notion of mixed-strategy equilibrium to the context of players whose preferences do not have expected utility representations; just as for the case of mixed-strategy equilibrium, it is assumed that each player explicitly randomizes to maximize his own utility given his correct belief about the strategy of the other. The definition of ambiguous act equilibrium differs from the definition of mixed-strategy equilibrium in that players are assumed to know the probability with which the different actions are chosen by the players according to the latter definition. While under mixed-strategy equilibrium, players need to pick objective mixtures as their strategies, subjectively mixed strategies are permitted in the case of ambiguous act equilibria. In Section 11.3, I discuss the option to use ambiguous games to define another equilibrium concept that would extend the notion of equilibrium in beliefs to the context of ambiguity-averse players. In analogy to the acronym AAE, I use the acronym *NE* to denote a Nash equilibrium p of a game $G = (A, \succ)$.

5. Alternative concepts of equilibrium

The definition of AAE proposed here differs from the definitions of equilibrium by Klibanoff (1996), Dow and Werlang (1994), Lo (1996), Eichberger and Kelsey (2000), Marinacci (2000) and Lehrer (2008) in that they all use some particular representation of preferences in their definitions of equilibrium. No particular representation of preferences is assumed in the present context.

Different from Dow and Werlang (1994), Eichberger and Kelsey (2000) and Marinacci (2000), who all extend the notion of equilibrium in beliefs (which I will discuss further in Section 11.3), the present paper assumes that players explicitly randomize. In that aspect, the definition advocated here is closer to the definitions presented by Klibanoff (1996), Lo (1996), and Lehrer (2008), as they all assume that players choose mixtures as their strategies. However, since Klibanoff (1996), Lo (1996), and Lehrer (2008) all assume that all players' strategy spaces contain only pure or mixed strategies, they cannot directly apply the Nash equilibrium concept to obtain an analogue to mixed-strategy equilibrium for the case of uncertainty-averse players. On the one hand, a player cannot know the other player's mixed strategy and at the same time be uncertain about it. Klibanoff (1996), Lo (1996), and Lehrer (2008) all assume that players maximize against their belief on the strategy of the other players; however, they do not assume that this belief *coincides* with the strategy of the other. Instead, their notions of equilibrium require some conditions of *consistency* between beliefs and strategies, which I will elaborate on further in Section 11.4.

This differs markedly from the present definition. I allow for ambiguous strategies. Consequently, the use of the Nash equilibrium concept to define AAE does not eliminate uncertainty. The novelty of the present approach lies in the definition of a game, the equilibrium concept itself is not new, I use Nash equilibrium. In case of the present definition, players maximize given their belief on the strategy of the other player, and this belief *coincides* with the actual strategy of the other. There does, however, remain some scope for uncertainty as no player needs to hold a probabilistic belief on the underlying state space.

The current paper is closest to the approaches proposed in Section 7 of Lo (1996), by Epstein (1997) and by Azrieli and Teper (2009). Lo constructs a state space and strategies mapping states to actions for all players to ground his equilibrium concept in an environment in which player preferences are defined over fundamentals. He acknowledges that the properties

of agreement and stochastic independence need to be defined for this environment and goes on to define these properties in terms of the particular representation of preferences he chose. Epstein embeds normal form games in a model of strategic uncertainty similar to mine. In contrast to my approach, he is interested in a notion of rationalizability for players whose preferences cannot be represented by expected utilities. He therefore does not impose any properties of strategic independence or common priors on the games he studies. While Azrieli and Teper's notion of games and equilibria with uncertainty-averse players is quite similar to mine, it differs insofar as neither correlation devices nor diverging beliefs are ruled out, and insofar as they allow for uncertainty about the environment as well as for uncertainty about other player's actions.

In Section 11.4, I will compare the main results of the present study with some of the results that are based on the above-mentioned definitions of equilibrium for uncertainty-averse players.

6. Observational equivalence: definition

The main claim of this study is that the Ambiguous Act Equilibria and the Nash equilibria of a two-player game $G = (A, \succsim)$ are observationally equivalent when the preferences of all players satisfy Schmeidler's uncertainty aversion in addition to some very standard assumptions on preferences, such as completeness and monotonicity. Observational equivalence captures the idea that an outsider who only observes the action profiles that players choose cannot tell whether the players are ambiguity-neutral or ambiguity-averse. In short, two strategy profiles are considered observationally equivalent if their supports coincide. To proceed further I need to define the notion of the support of an ambiguous act.

Definition 3. An action profile a is said to be in the *support* of strategy profile f if there exists a non-null state ω such that $f(\omega)(a) > 0$. The set of all actions in the support of f by is denoted by $\text{supp}(f)$. Two acts f, g are called *observationally equivalent* if they have the same support. Two sets of acts \mathcal{F}, \mathcal{G} are called observationally equivalent if for every $f \in \mathcal{F}$ there is an observationally equivalent $g \in \mathcal{G}$ and vice versa.

The support of a constant act p equals the support of the lottery p in the usual sense. It is important to note that players agree on null-events in ambiguous games. The present notion of the support of a strategy profile cannot be applied to general ambiguous games as the set of 1-null states need not coincide with the set of 2-null states in such a game. Without this assumption the notion of "support" is not well defined.

The question underlying the definition of observational equivalence is: is there any action profile that is consistent with equilibrium play among players with a particular attitude towards ambiguity – neutral or averse – without being consistent with equilibrium play among players with a different attitude towards ambiguity? In other words, is there any action profile that "proves" that players are ambiguity-averse in the sense that this action profile is contained in the support of ambiguous act equilibrium, but is not contained in the support of any mixed-strategy equilibrium?

So one might ask: can an observer not use more than just *one* action profile to determine whether the players are ambiguity-neutral or -averse? Aren't the frequencies with which all action profiles are being played also observable? Yes, they are; however, the question whether some frequencies are consistent with a subjective act is very much open to debate. To see this, take an act $f : \Omega \rightarrow \Delta(A)$ with $\Omega = \{\omega_1, \omega_2\}$, $A = \{a_1, a_2\}$ and $f(\omega_1) = a_1$, $f(\omega_2) = a_2$ and both ω_1 and ω_2 non-null. This act entails no prediction about the frequency of the occurrence of a_1 and a_2 . I would need to impose further assumptions on players' preferences to relate the observed frequencies to the played acts f . I chose to avoid this by using the equality of support as my criterion of observational equivalence. However, ambiguity aversion could be used to explain frequencies that cannot be produced by mixed strategy equilibria as equilibrium phenomena. As long as players' learning is too limited to learn the correct frequencies, many different frequencies are consistent with the predictions of ambiguous act equilibria. Keeping this in mind, the theory of ambiguous act equilibria can be used to interpret some of the existing empirical evidence against mixed strategy as equilibrium play of ambiguity-averse players.

7. Preferences

Until now I have not specified the preferences of the players beyond requiring the properties of independent strategies and agreement on non-null events. To get any results some further requirements will have to be imposed. The four basic properties defined next will be assumed from here on, if not explicitly mentioned otherwise:

(RAT) Preferences are transitive and complete.

(EU) Preferences over constant acts – that is, preferences over lotteries – have an expected utility representation; for any player i there exists an affine function $u_i : \Delta(A) \rightarrow \mathbb{R}$ that represents the players preferences over constant acts (lotteries) $\Delta(A)$.⁴

⁴ Clearly, I could have stated some more basic properties on the player's preferences over constant acts that imply (EU). I chose to state these assumptions summarily as "(EU)" for the sake of brevity.

(MON) Take two acts f, g . If $f(\omega) \succsim_i g(\omega)$ holds for all non-null states $\omega \in \Omega$, then $f \succsim_i g$. If, in addition, $f(\omega') \succ_i g(\omega')$ holds for some non-null state ω' , then $f \succ_i g$.

(UA) For any two acts $f, g : \Omega \rightarrow \Delta$ with $f \sim g$, it holds that $\alpha f + (1 - \alpha)g \succsim f$.

The first two assumptions (RAT) and (EU) are very much standard, I only mention in passing that Bewley's (2002) representation of preferences for Knightian uncertainty does not assume completeness, and therefore violates (RAT) as defined here. The definition of monotonicity provided here is somewhat stronger than the most commonly used definition of monotonicity (MON') which only requires that $f \succsim_i g$ holds for two acts f, g with $f(\omega) \succsim_i g(\omega)$ for all non-null states $\omega \in \Omega$. To see the difference between the two versions of monotonicity, observe that the stronger axiom (MON) requires that an act f is strictly preferred to an act g , if f is strictly preferred to g for a non-null set of states ($f(\omega) \succ_i g(\omega)$ for $\omega \in E$ with E non-null) and indifferent on all other states ($f(\omega) \sim_i g(\omega)$ for $\omega \notin E$). In this case, the weaker axiom (MON') only requires that f is weakly preferred to g . Following the proof of the main result of this paper, Theorem 1, I will discuss how the result would change if one was to replace (MON) by the weaker (MON'). The axiom (UA) was introduced by Schmeidler (1989). He defined “ambiguity aversion” as a preference for randomization: if an agent is indifferent between two uncertain acts, then he should like an objective mixture of these two acts at least as much as either one of them.

Preferences that can be represented by an expected utility satisfy all of the four axioms. Gilboa and Schmeidler (1989) show that preferences satisfy (RAT), (EU), (MON') and (UA) in addition to a continuity requirement and a restriction of the independence axiom to mixtures with constant acts, if and only if preferences have a *maxmin expected utility representation* $U(f) = \min_{\pi \in C} \sum_{\omega \in \Omega} \pi(\omega)u(f(\omega))$ for some convex and compact set of priors C and an expected utility functional $u : \Delta(A) \rightarrow \mathbb{R}$. The stronger monotonicity axiom (MON) holds for a maxmin expected utility representation, if and only if all priors π in the set C are mutually absolutely continuous, in the sense that $\pi(E) > 0$ for some $\pi \in C$ holds if and only if $\pi'(E) > 0$ for all $\pi' \in C$ and all $E \subset \Omega$. Mutual absolute continuity of the priors in the maxmin expected utility representation has been axiomatized by Klibanoff (1996) and Epstein and Marinacci (2007). Preferences that have a Choquet expected utility representation following Schmeidler (1989) always satisfy (RAT), (EU) and (MON'), they satisfy (UA), if and only if the capacity is convex. Cerreia-Vioglio et al. (2008) characterize the representation of all preferences that satisfy (RAT), (EU), (UA), (MON') and some form of continuity. Variational Preferences following Maccheroni et al. (2006), Multiplier Preferences following Hansen and Sargent (2001) and Smooth Ambiguity-Averse Preferences following Klibanoff et al. (2005) all constitute special cases of the preferences described in Cerreia-Vioglio et al. (2008).⁵ Abstracting from the slight strengthening of the monotonicity axiom the four axioms given here (RAT), (EU), (MON) and (UA) encompass a very large set of preferences described in the literature.

8. Observational equivalence: the main result

The assumption of (RAT), (EU), (MON) and (UA) is sufficient to show that the set of Ambiguous Act Equilibria of a game is observationally equivalent to the set of its Nash equilibria. If these 4 axioms are valid, then there does not exist a single action profile that is in the support of an equilibrium when players are ambiguity-averse without being in the support of any Nash equilibrium of the game or vice versa. This is the main result of this paper.

Theorem 1. *Let $G = (A, \succsim)$. The set of AAE of G is observationally equivalent to the set of NE of G .*

The next section is devoted to showing the preliminary result that a constant act p is a NE of G , if and only if it is an AAE of G . This already allows us to conclude that there are no NE that are not observationally equivalent to any AAE. The following section establishes the converse conclusion.

9. Mixed strategy equilibria

In this section, it is shown that a strategy profile p is an NE of a game $G = (A, \succsim)$ if and only if it is an AAE of that game. The main building block towards that result is the observation that for any strategy profile f , player 1 has a mixed strategy p_1^* that is weakly preferred to f_1 , keeping the other player's strategy fixed at f_2 . This observation is the content of the next lemma.

Lemma 2. *Let $G' = (\Omega, A, \succsim')$. For any strategy profile f and either player $i = 1, 2$ there exists a mixed strategy p_i^* such that $(p_i^*, f_{-i}) \succsim_i f$ and $\text{supp}(p_i^*) = \text{supp}(f_i)$.*

Proof. If there exists a p_i such that the event $\{\omega_i : f_i(\omega_i) \neq p_i\}$ is null, then $(p_i, f_{-i}) \sim_i f$ and $\text{supp}(p_i) = \text{supp}(f_i)$ both hold and the proof is complete. If no such event exists, f_i can be represented as $f_i = (p_i : E, q_i : F, f_i : \overline{E \cup F})$ with $p_i \neq q_i$ and $E, F \subset \Omega_i$ both non-null. Since Ω_1 and Ω_2 are stochastically independent an α can be found such that either

⁵ The axiom (MON) also corresponds to an assumption of mutual absolute continuity in the case of the Cerreia-Vioglio et al. (2008) framework (and all its special cases). I owe this insight to a private conversation with Massimo Marinacci. A proof is available from the author upon request.

$$f = ((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i}) \sim'_k ((\alpha p_i + (1 - \alpha)q_i : E, p_i : F, f_i : \overline{E \cup F}), f_{-i})$$

or

$$f = ((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i}) \sim'_k ((q_i : E, \alpha p_i + (1 - \alpha)q_i : F, f_i : \overline{E \cup F}), f_{-i}).$$

Assume w.l.o.g. that the first indifference holds. Uncertainty aversion implies that i prefers any mixture of the two indifferent acts to f , so we have in particular that

$$\begin{aligned} f &= ((p_i : E, q_i : F, f_i : \overline{E \cup F}), f_{-i}) \\ &\succ'_i ((\beta p_i + (1 - \beta)q_i : E, \beta p_i + (1 - \beta)q_i : F, f_i : \overline{E \cup F}), f_{-i}) \\ &= ((\beta p_i + (1 - \beta)q_i : E \cup F, f_i : \overline{E \cup F}), f_{-i}) \end{aligned}$$

for $\beta = \frac{1}{2-\alpha}$. If $\Omega_i^m \subset E \cup F$ the proof is complete. If not, the same step can be repeated once again. The finiteness of the state space Ω_i implies that this process will end after finitely many repetitions of this step: a p_i^* such that $(p_i^*, f_2) \succ_i f$ can be found.

To see that $\text{supp}(p_i^*) = \text{supp}(f_i)$, observe first of all that $\text{supp}(p_i^*) \subset \text{supp}(f_i)$ as p_i^* is a mixture over the lotteries $f_i(\omega_i)$ for all non-null $\omega \in \Omega$. Next, observe that the mixing parameter β constructed above assigns positive probability to both p_i and q_i ; the same would have held for the alternative case in which we would have found an α such that $f \sim_i ((q_i : E, \alpha p_i + (1 - \alpha)q_i : F, f_i : \overline{E \cup F}), f_{-i})$. We conclude that the lottery over lotteries $f_i(\omega_i)$ has full support and therefore $\text{supp}(p_i^*) = \text{supp}(f_i)$. \square

The assumption of independent state spaces Ω_1 and Ω_2 is crucial for the above proof as it permits the generation of indifferent acts by permuting the strategies played in different events. Uncertainty aversion implies that any convex combination of these indifferent acts is at least weakly preferred. Finally, it is a matter of simple algebra to show that there always exists such a mixture that evens out any uncertainty over the different events that were involved in the permutation of strategies.

Theorem 2. Let $G = (A, \succ)$. A mixed-strategy profile p is an AAE of G if and only if p is an NE of G .

Proof. Let p be an AAE of the ambiguous act extension $G' = (\Omega, A, \succ')$ of G . So $p \succ'_i (f_i, p_{-i})$ holds for all deviations f_i in particular it holds for deviations p'_i , so $p \succ_i (p'_i, p_{-i})$ holds and therefore p is an NE of G' . Next, assume that p is an NE of G' . Suppose p was not an AAE of G , that is suppose that there exists a deviation f_i for player i such that $(f_i, p_{-i}) \succ'_i p$. By Lemma 2 we have that there exists a p_i^* such that $(p_i^*, p_{-i}) \succ'_i (f_i, p_{-i}) \succ'_i p$ a contradiction with the assumption that p was a Nash equilibrium. \square

Theorem 2 should not come as a big surprise. As shown in the preceding lemma, uncertainty aversion implies that players can always be made weakly better off by objectively mixing over all the lotteries they use according to their ambiguous act strategies. Given that the opponent plays a mixed strategy p_{-i} , player i cannot be made strictly better off by an ambiguous strategy. It should be noted from the proof though that Theorem 2 could be strengthened to say that p is an NE of G , if and only if it is an AAE of any ambiguous act extension of G . As a direct consequence we get the following corollary.

Corollary 1. Any ambiguous game $G = (\Omega, A, \succ)$ has an AAE.

10. Observational equivalence: the proof

To prove Theorem 1, I first show that mixtures that are “dominated”, in a sense to be defined precisely, will never be used in an AAE (Lemma 3). I will then go on to show that this implies that a player’s belief on the equilibrium strategy of the other can always be represented by a probability (Lemma 4). These two lemmata yield the proof of the main result of this paper (Theorem 1). It is convenient to use matrix algebra to state and prove all these results. Some more notation needs to be introduced.

A generic vector x is assumed to be a column vector, row vectors are obtained by taking the transpose x^T . For any two vectors x, y of the same length, the relations “ $>$ ” and “ \geq ” are defined by $x \geq y$, if and only if $x_t \geq y_t$ for all components t and $x > y$, if and only if $x \geq y$, but not $y \geq x$. For any $\alpha \in \mathbb{R}$, I denote the vector $(\alpha, \alpha, \dots, \alpha)^T$ by $\bar{\alpha}$.

For a fixed strategy profile f , define an $|A_1| \times |\Omega_2^m|$ -matrix $X(f)$ with $u_1(a_1^j, f_2(\omega_2^k))$, the entry in row j and column k , the expected utility player 1 derives from playing a_1^j when player 2 plays the mixed strategy $f_2(\omega_2^k)$. The submatrix of X that consists only of the rows j associated with actions $a_1^j \in \text{supp}(f_1)$ is called $Y(f)$; the complementary submatrix (consisting only of the rows associated with actions that 1 never plays under f) is called $Z(f)$. If $p_1^T X(f) > q_1^T X(f)$ holds for some $p_1, q_1 \in \Delta(A_1)$, then we say that p_1 dominates q_1 at f_2 . Observe that $p_1^T X(f) > q_1^T X(f)$ implies that player 1

weakly prefers playing p_1 to playing q_1 in every non-null state ω_2 and sometimes strictly prefers doing so. Consequently, given the assumption of (MON), we have that $(p_1, f_2) \succ_1 (q_1, f_2)$ holds, when p_1 dominates q_1 at f_2 . Since I consider a fixed strategy profile f , I drop the argument f in the expressions $X(f)$, $Y(f)$ and $Z(f)$ in the sequel.

Lemma 3. Let $G' = (\Omega, A, \succ')$. Let f be an AAE of G' . Then, there do not exist any $p_1, q_1 \in \Delta(A_1)$ with $\text{supp}(q_1) \subset \text{supp}(f_1)$ such that p_1 dominates q_1 at f_2 .

Proof. Suppose there existed some $p_1, q_1 \in \Delta(A_1)$ with $\text{supp}(q_1) \subset \text{supp}(f_1)$ such that and $p_1^\top X > q_1^\top X$. By Lemma 2, there exists a mixed strategy $p_1^* \in \Delta(A_1)$ such that $(p_1^*, f_2) \succ'_1 f$ and $\text{supp}(p_1^*) = \text{supp}(f_1)$. Since $\text{supp}(p_1^*) = \text{supp}(f_1) \supset \text{supp}(q_1)$ there exists some $r_1 \in \Delta(A_1)$ with $\text{supp}(r_1) \subset \text{supp}(f_1)$ and an $\gamma \in (0, 1]$ such that $p_1^* = \gamma q_1 + (1 - \gamma)r_1$. We obtain that

$$(p_1^*)^\top X = (\gamma q_1 + (1 - \gamma)r_1)^\top X < (\gamma p_1 + (1 - \gamma)r_1)^\top X.$$

So $\gamma p_1 + (1 - \gamma)r_1$ dominates p_1^* at f_2 which implies that $((\gamma p_1 + (1 - \gamma)r_1), f_2) \succ'_1 (p_1^*, f_2) \succ'_1 f$, a contradiction with the assumption that f is an AAE. \square

The next lemma shows that it is precisely the condition derived in the prior lemma, namely the absence of two $p_1, q_1 \in \Delta(A_1)$ with $\text{supp}(q_1) \subset \text{supp}(f_1)$ such that p_1 dominates q_1 at f_2 , that allows us to conclude that there exists a \tilde{p}_2 with $\text{supp}(\tilde{p}_2) = \text{supp}(f_2)$ such that player 1 is indifferent between playing any action $a_1^j \in \text{supp}(f_1)$ when player 2 plays \tilde{p}_2 and (weakly) prefers playing any of these actions to playing any of the other actions $a_1^j \notin \text{supp}(f_1)$. The mixed strategy \tilde{p}_2 is constructed as the composite lottery of some lottery p_2 over all mixed strategies $f_2(\omega_2)$ that player 2 plays according to f_2 . Using matrix algebra, the desired result can be expressed as follows: there exists an $\alpha \in \mathbb{R}$ and a $p_2 \in \Delta(\Omega_2^{nm})$ such that $Yp_2 = \bar{\alpha}$ and $Zp_2 \leq \bar{\alpha}$, where the matrices Y and Z are defined as described above.

Lemma 4. Let $G' = (\Omega, A, \succ')$. Suppose there do not exist any $p_1, q_1 \in \Delta(A_1)$ with $\text{supp}(q_1) \subset \text{supp}(f_1)$ such that p_1 dominates q_1 at f_2 . Then, there exists a probability $p_2 \in \Delta(\Omega_2^{nm})$ with full support and an $\alpha \in \mathbb{R}$ such that $Yp_2 = \bar{\alpha}$ and $Zp_2 \leq \bar{\alpha}$.

Proof. (\Rightarrow) Suppose there exists a $p_2 \in \Delta(\Omega_2^{nm})$ with full support and an $\alpha \in \mathbb{R}$ such that $Yp_2 = \bar{\alpha}$ and $Zp_2 \leq \bar{\alpha}$. Suppose we also had $p_1, q_1 \in \Delta(A_1)$ with $\text{supp}(q_1) \subset \text{supp}(f_1)$ such that $p_1^\top X > q_1^\top X$. This yields a contradiction as $\alpha = q_1^\top \bar{\alpha} = q_1^\top Xp_2 < p_1^\top Xp_2 \leq p_1^\top \bar{\alpha} = \alpha$.

(\Leftarrow) Suppose there exists no $p_2 \in \Delta(\Omega_2^{nm})$ with full support and $\alpha \in \mathbb{R}$ such that $Yp_2 = \bar{\alpha}$ and $Zp_2 \leq \bar{\alpha}$. This is equivalent to $S \cap r = \emptyset$ for $S, r \subset \mathbb{R}^{|A_1|}$, $r := \{\bar{\alpha} \mid \alpha \in \mathbb{R}\}$, and

$$S := \{s \mid s_Y = Yp_2 \text{ and } s_Z \geq Zp_2 \text{ for some } p_2 \in \Delta(\Omega_2^{nm}) \text{ with } \text{supp}(p_2) = \Omega_2^{nm}\}$$

where $s_Y = (s_j)_{a_1^j \in \text{supp}(f_1)}$ and $s_Z = (s_j)_{a_1^j \notin \text{supp}(f_1)}$ are defined as the subvectors of s that correspond to actions a_1^j in the support and outside the support of f_1 respectively. Since S is a convex set, there exists a separating hyperplane H such that $r \subset H$ and $H \cap S = \emptyset$. Let this plane H be described by a vector λ such that $\lambda^\top x = 0$ implies $x \in H$ and $\lambda^\top x > 0$ for all $x \in S$. Since $\bar{1} \in r \subset H$ we have that $\sum \lambda_i = 0$.

Next define two vectors κ and ρ by $\kappa_l = \lambda_l$ if $\lambda_l > 0$, $\kappa_l = 0$ otherwise, $\rho_l = -\lambda_l$ if $\lambda_l < 0$ and $\rho_l = 0$ otherwise. Observe that $\sum \kappa_l = \sum \rho_l > 0$.⁶ Define $\tilde{\lambda}, \tilde{\kappa}$ and $\tilde{\rho}$ by

$$\tilde{\lambda}_l = \frac{\lambda_l}{\sum \kappa_l}, \quad \tilde{\kappa}_l = \frac{\kappa_l}{\sum \kappa_l}, \quad \tilde{\rho}_l = \frac{\rho_l}{\sum \kappa_l}.$$

Observe that $\tilde{\lambda}$ and $\tilde{\lambda}$ as normal vectors describe the same plane and that $\tilde{\rho}, \tilde{\kappa}$ are elements of $\Delta(A_1)$. Consequently we have that $\tilde{\lambda}^\top x > 0$ for all $x \in S$. As $\tilde{\lambda} = \tilde{\kappa} - \tilde{\rho}$ we have that $\tilde{\kappa}^\top x > \tilde{\rho}^\top x$ for all $x \in S$.

I show next that $\tilde{\rho}_j = 0$ for all j that are associated with actions $a_1^j \notin \text{supp}(f_1)$. Suppose we had $\tilde{\rho}_j > 0$ for such a j . Fix an $x \in S$, observe that $\tilde{\kappa}^\top x > \tilde{\rho}^\top x$ has to hold for this x , as this has to hold for all $x \in S$. Next define \tilde{x} by $\tilde{x}_{-j} = x_{-j}$ and $\tilde{x}_j > \frac{\tilde{\kappa}^\top x - \tilde{\rho}_j^\top x_{-j}}{\tilde{\rho}_j}$. By our construction of S we can find such an \tilde{x} that is also an element of S . Observe that

$$\tilde{\rho}^\top \tilde{x} = \tilde{\rho}_{-j}^\top \tilde{x}_{-j} + \tilde{\rho}_j \tilde{x}_j > \tilde{\rho}_{-j}^\top x_{-j} + \tilde{\rho}_j \frac{\tilde{\kappa}^\top x - \tilde{\rho}_{-j}^\top x_{-j}}{\tilde{\rho}_j} = \tilde{\kappa}^\top x = \tilde{\kappa}^\top \tilde{x}$$

where the very last equality follows from the fact that, on the one hand, $\tilde{\rho}_j > 0$ implies $\tilde{\kappa}_j = 0$ and, on the other hand $x_{-j} = \tilde{x}_{-j}$. But $\tilde{\rho}^\top \tilde{x} > \tilde{\kappa}^\top \tilde{x}$ stands in contradiction with $\tilde{\kappa}^\top x > \tilde{\rho}^\top x$ holding for all $x \in S$. We conclude that $\tilde{\rho}_j = 0$ for all j that are associated with actions $a_1^j \notin \text{supp}(f_1)$. Therefore $\text{supp}(\rho)$ is a subset of $\text{supp}(f_1)$.

⁶ The vectors κ and ρ are defined such that $\sum \kappa_l = \sum \rho_l \geq 0$. If we had that $\sum \kappa_l = \sum \rho_l = 0$, we would also have $\lambda = 0$, a contradiction with the assumption that λ describes the hyperplane H .

To conclude this proof, observe that $\tilde{\kappa}X^k \geq \tilde{\rho}X^k$ for all columns k of the matrix X as any such column X^k can be approached by a sequence $x_n \in S$. Finally, it cannot be true that $\tilde{\kappa}U^k = \tilde{\rho}X^k$ for all columns k , as we could then find $x \in S$ with $\tilde{\kappa}x = \tilde{\rho}x$. So it must be true that $\tilde{\kappa}X^{k'} > \tilde{\rho}X^{k'}$ for some columns k' . Hence, we found two probabilities $\tilde{\kappa}$ and $\tilde{\rho}$ such that $\tilde{\kappa}^T X > \tilde{\rho}^T Y$. \square

Combining the preceding two lemmata yields the main result of this paper, which is proved next:

Proof of Theorem 1. (\Leftarrow) Let p be an NE of G , then by Theorem 2 p itself is an AAE of G , so G has an AAE with the same support.

(\Rightarrow) Let f be an AAE of G . Following Lemma 3, there do not exist any $p_1, q_1 \in \Delta(A_1)$ with $\text{supp}(q_1) \subset \text{supp}(f_1)$ such that p_1 dominates q_1 at f_2 . Applying Lemma 4, we conclude that there exists a probability p_2 on Ω_2^m with full support such that all $a_1^j \in \text{supp}(f_1)$ are best replies to \tilde{p}_2 which is defined by $\tilde{p}_2(a_2) = \sum_{\omega_2 \in \Omega_2} p_2(\omega_2) f(\omega_2)(a_2)$. Construct \tilde{p}_1 in the same fashion exchanging the names of players 1 and 2 in the above results. The mixed-strategy profile \tilde{p} is an NE of G with $\text{supp}(\tilde{p}) = \text{supp}(f)$, as all $a_i \in \text{supp}(f_i)$ are best replies to \tilde{p}_{-i} for $i = 1, 2$. \square

Theorem 1 is the main result of this paper. This result establishes that an outside observer cannot distinguish the behavior of two uncertainty-averse players, from the behavior of two uncertainty-neutral players when he observes only the outcomes of their play. Of course, certain conditions have to hold for this result to apply: it is shown that observational equivalence holds for two-player games, where both players are expected utility maximizers with respect to lotteries, have monotonic preferences, and satisfy Schmeidler’s axiom of uncertainty aversion. The assumption “(EU)” was used to define the matrix U , the assumption “(MON)” was used to conclude that $(p_1, f_2) \succ'_1 (q_1, f_2)$ holds when p_1 dominates q_1 at f_2 , and finally “(UA)” was used in Lemma 2, which in turn played an important role in the proof of Lemma 3.

11. Discussion

11.1. Monotonicity

Weakening (MON) to (MON') the domination of q_1 by p_1 at f_2 only implies that $(p_1, f_2) \succ'_1 (q_1, f_2)$. To see that (MON') is not sufficient to obtain Theorem 1, consider the following example of a game that has an AAE that is not observationally equivalent to any NE.

Example 1. Let the ambiguous game $G' = (\Omega, A, \succ')$ be defined by Ω_1 a singleton, $\Omega_2 = \{\omega_1, \omega_2\}$ and the payoff matrix:

	L	R
T	10, 1	0, 0
B	11, 0	0, 1

Assume that both players have a maxmin expected utility representation following Gilboa and Schmeidler (1989) with $C_1 = [0, 1/2]$ and $C_2 = \{1/2\}$ the players’ sets of beliefs on ω_1 . Having maxmin expected utility representations both players’ preferences automatically satisfy (MON'). On the other hand, player 1’s preferences do not satisfy the stronger condition (MON) as the priors in C_1 are not mutually absolutely continuous. Observe that the strategy profile f according to which player 1 equally randomizes over his two actions and $f_2(\omega_1) = L$ and $f_2(\omega_2) = R$ has full support and is an AAE of the game. To see this, it is important to note that, given the belief set C_1 and player 2’s strategy f_2 , player 1 completely discounts the occurrence of the state ω_1 which is more favorable for him. However, ω_1 is not a null state, for some different payoffs ω_1 could very well be relevant for player 1’s preference over two acts that differ only on ω_1 . The game does not have an NE with full support.

The equilibrium constructed in the prior example strikes me as particularly unappealing. Why would player 1 choose to play T with positive probability, when playing B is never worse and strictly better in some non-null event? Since the equilibrium predictions using (MON') differ from the equilibrium predictions using the stronger axiom (MON) only in that under the more common axiom of monotonicity players might use strategies in equilibrium that are never better than some others and that can be improved upon strictly for some non-null events, I chose to study mainly preferences satisfying (MON).

11.2. Games with more than two players

Most definitions in this paper can easily be extended to n -player games. In the present section, I will show that the main result of this paper, Theorem 1, does not extend to n -player games. A theory of games with more than two ambiguity-averse players carries the potential to yield substantially different predictions from the standard theory of mixed-strategy equilibrium. I will first discuss an example of a 3-player game with an AAE that is observationally different from any NE

in that game. I will then go on to provide some reasons why a detailed study of this question lies beyond the scope of this paper, as the basic understanding of “common priors” developed here does not suffice to tackle the case of n -players. A better grasp of this concept is needed to fully understand the case of games with more than two players. The following example builds on Example 2.3 in Aumann (1974).

Example 2. Consider the three player game given by the following payoff matrices:

	l	
	L	R
T	0, 8, 0	3, 3, 3
B	1, 1, 1	0, 0, 0

	r	
	L	R
T	0, 0, 0	3, 3, 3
B	1, 1, 1	8, 0, 0

Let Ω_1 and Ω_2 be singletons and let $\Omega_3 = \{\omega_1, \omega_2\}$; assume furthermore that all three players’ preferences can be represented by maxmin expected utility functions with $C_i = [\frac{1}{4}, \frac{3}{4}]$ for $i = 1, 2, 3$.

The strategy profile f with $f_1 = T, f_2 = R$ and $f_3(\omega_1) = l, f_3(\omega_2) = r$ is an AAE of this game. To see this, observe that player 3 does not have an incentive to deviate as his utility is 3 no matter which of the two boxes he picks. Secondly players 1 and 2’s utilities from playing their pure strategies against f_{-i} the fixed strategies of the others can be calculated as:

$$U_1(T, f_{-1}) = \min_{\pi \in [\frac{1}{4}, \frac{3}{4}]} \pi \times 3 + (1 - \pi) \times 3 = 3,$$

$$U_1(B, f_{-1}) = \min_{\pi \in [\frac{1}{4}, \frac{3}{4}]} \pi \times 0 + (1 - \pi) \times 8 = \frac{3}{4} \times 0 + \frac{1}{4} \times 8 = 2 < 3,$$

$$U_2(R, f_{-2}) = \min_{\pi \in [\frac{1}{4}, \frac{3}{4}]} \pi \times 3 + (1 - \pi) \times 3 = 3,$$

$$U_2(L, f_{-2}) = \min_{\pi \in [\frac{1}{4}, \frac{3}{4}]} \pi \times 8 + (1 - \pi) \times 0 = \frac{1}{4} \times 8 + \frac{3}{4} \times 0 = 2 < 3.$$

Aumann (1974) shows that the game in Example 2 has no NE with (TRr) or (TRL) in its support. At the same time, he shows that the game has an equilibrium with non-common priors in which the first two players play T, R . A necessary condition for the existence of such a “correlated equilibrium” is that player 1 assigns a higher probability to player 3 playing l than does player 2. Aumann shows, in particular, that there is such an equilibrium if player 1 assigns a probability of $3/4$ to player 3 choosing l , whereas player 2 only assigns a probability of $1/4$ to that event. The set $C_i = [\frac{1}{4}, \frac{3}{4}]$ used in the example, implies that the two first players will use exactly these priors when calculating their respective *minimal* expected utility of strategy profile f in the above example.

So is a game theory with ambiguity-averse players going to herald a revival of game theory without common priors? Is any NE without common priors observationally equivalent to an AAE with ambiguity-averse players? Yes it is – in a very unsatisfactory way. To see this, observe that ambiguous games with expected utility maximizers allow for non-common priors as long as all players agree on zero-probability events. An equivalence result between the AAE and the NE with non-common priors would not be driven by the players ambiguity aversion (in fact, players would be assumed to be expected utility maximizers). The assumption that players only agree on null-events is strong enough for the purpose of the present paper: the observational equivalence result for two players could be derived using only this weak assumption on the players’ agreement of beliefs. The condition would have to be strengthened considerably for a study of games with more than two players.

It could be interesting to derive a non-probabilistic axiom of “common beliefs” from the observation that for the case of risk-averse expected utility maximizers, the common prior assumption holds, if and only if full insurance allocations are Pareto optimal. Billot et al. (2000) use exactly this observation to develop a notion of “sharing beliefs” for the case of maxmin expected utility maximizers. Example 2 withstands such a strengthening of the requirement of “common beliefs”. In fact, the above example would probably satisfy many different notions of “common beliefs” for uncertainty-averse agents as the sets of beliefs of all three players exactly *coincide* in the example.⁷

11.3. Equilibrium in beliefs

For the case of expected utility maximizers, any mixed-strategy equilibrium can be interpreted as an equilibrium in beliefs. According to this interpretation it is assumed that each player picks a pure strategy. Players do not know the pure

⁷ Lo (1996) covers games with more than 2 players. He assumes that preferences can be represented following Gilboa and Schmeidler (1989) and assumes that the belief sets of all players have to be equal. This is unsatisfactory as we do not have an axiomatic foundation for the assumption that belief sets should coincide exactly. Eichberger and Kelsey (2000) acknowledge that common priors matter for the context of games with more than 2 players. They do not attempt to tackle these questions in their article.

strategies picked by their opponents; instead they form (probabilistic) beliefs about the actual choices of their opponents. A set of such beliefs constitutes an equilibrium in beliefs if every action in the support of a player's belief is a best response given the belief of the player choosing that action. For the case of a two-player game with expected utility maximizers the distinction between equilibrium in beliefs and mixed-strategy equilibrium is purely philosophical. Two probabilities on the set of all actions of the two players constitute a mixed-strategy equilibrium, if and only if they constitute an equilibrium in beliefs.

This changes dramatically once we allow for ambiguously mixed strategies and ambiguous beliefs. To see why this would be true, observe that objective mixing allows an agent to smooth out subjective uncertainty. So some actions that are not best responses when played as pure strategies can be played with positive probability in an AAE. To formally state and prove the divergence of the two notions of equilibrium, observe first of all that the definition of AAE presented in the present paper generalizes the standard definition of mixed-strategy equilibrium in the sense that it allows for ambiguously mixed strategies in addition to objectively mixed strategies.

Just as for the definition of AAE, we take ambiguous games as the starting point for the definition of equilibria in beliefs. An act f_i is no longer interpreted as the strategy chosen by player i , but rather as player $-i$'s belief about the pure strategy of player i . This generalizes the standard definition of equilibrium in beliefs insofar as the belief on the strategy of the other need not be an objective mixture, but can be any act. Consequently, we could call a strategy profile f an equilibrium in beliefs, if every action in the support of f_i is a best response (among the pure strategies) for player i given his belief f_{-i} about the other player's strategy.

Some gaps need to be filled in this definition. Just as for the case of AAE, the difference between equilibria in ambiguous beliefs and equilibria in beliefs should not be driven by any extraneous changes in the definition of a game. In short, we need to check whether the conditions of strategic independence and agreement provided above can also be used in the present context. The condition that players agree on non-null events would amount to some condition of "non-delusion". Just as in the above case, the condition allows us to construct the notion of the "support" of an act. The definition of stochastic independence needs to differ though. Stochastic independence in the present context should mean that player i 's ranking over his response to a belief f_{-i} about the other should not depend on the event ω_j , which is observed by player i , when the state is ω . In the context of expected utility-maximizing agents this definition of stochastic independence is equivalent to the definition of stochastic independence given above. Outside of this context, these definitions need not coincide. With all this in mind, we can now give an example of a game in which equilibria in beliefs and equilibria in ambiguous beliefs do not coincide.

Example 3. Let the ambiguous game $G' = (\Omega, A, \succ')$ be defined by a singleton Ω_1 and $\Omega_2 = \{\omega_1, \omega_2\}$ and the following payoff matrix:

	L	R
T	0, 0	1, 0
M	1, 0	0, 0
B	0.4, 0	0.4, 0

Observe that this example represents a version of the classical Ellsberg example as a game between the experimental subject (player 1) and the experimenter (player 2). The experimenter is not supposed to have any preferences over the outcome of the game, so his utility is evaluated at the constant value 0. Now interpret the two states of Ω_2 of drawing a blue or a green ball from a classical two-color Ellsberg urn. If the subject is uncertainty-averse, he might prefer to obtain 0.4 for sure rather than obtaining 1 only if a ball of a particular color is drawn from the urn. In that case the act $f_1 = B$ and $f_2(\omega_1) = L, f_2(\omega_2) = R$ is an equilibrium in ambiguous beliefs. This strategy profile is not an AAE, as the subject would be strictly better off by objectively mixing between T and M .

The marked difference between the two solution concepts generates the need to take sides. Which one of the two concepts is the appropriate solution concept? The answer very much depends on whether players (or people) are able to commit to randomize their strategies. Would a subject in the above experiment be able to commit to a mixed strategy? Would she be able to ignore her ability to reconsider the action prescribed by the mixture? Answers to questions like this one certainly depend on (a) the time frame of the decisions and (b) what is at stake. Either extreme, namely the assumption that players either are able to commit fully as well as the assumption that players are completely incapable to commit seems too strong. In the context of expected utility-maximizers, the two assumptions yield the same predictions, so there was no need to generate an intermediate assumption. In the present context of ambiguity-averse players, the predictions strongly depend on the assumption on the ability to commit to randomized strategies. Further research may find a good way to compromise between the two extreme assumptions.

11.4. Observational difference: other equilibrium concepts

I already pointed out in Section 5 that prior definitions build on particular representations of the players preferences, whereas the present definition does not make use of a particular representation. Klibanoff (1996) as well as Lo (1996)

use Gilboa and Schmeidler's maxmin expected utilities to define their respective equilibrium concepts, whereas Dow and Werlang (1994), Eichberger and Kelsey (2000) as well as Marinacci's (2000) use Schmeidler's Choquet utility representation to define their equilibrium concepts. It is puzzling then, that only Lo finds that "the observed behavior of the uncertainty-averse players (the actual strategies they choose) is consistent with [expected] utility maximization" (Lo, 1996, p. 463), even though the preferences assumed by all the mentioned studies satisfy (RAT), (EU), (MON') and (UA). I already showed above in Section 11.1, that the weakening of (MON) to (MON') allows for AAE that are observationally different from all NE of a game. To see that the difference between the results of the preceding studies and Theorem 1 depend on more than just the difference in the monotonicity axiom, let me discuss Klibanoff's notion of equilibrium at the hand of an example. A mixed-strategy profile is an equilibrium following Klibanoff if each player maximizes their utility given their belief on the other and if a consistency condition between beliefs and actually chosen strategies holds. The consistency condition requires that the (mixed) strategy of player i must be part of player $-i$'s belief set C_i about the strategy of player i . Formally the definition can be stated as follows:

Definition 4. Take a game $G = (A, \succ)$. A profile of mixed strategies p^* is considered an *equilibrium* or short *KE*, if there exist two convex and compact belief sets $C_i \subset \mathcal{P}(A_i)$, $i = 1, 2$, such that p_i^* maximizes $\min_{p_{-i} \in C_{-i}} u^i(p_i \times p_{-i})$ and $p_i^* \in C_i$ for $i = 1, 2$.

Let me use the following example of Klibanoff (1996) to show that the difference between KE and AAE goes beyond the assumption of different axioms of monotonicity.

Example 4. Let the normal form game $G = (A, \succ)$ be defined by the payoff matrix:

	L	R
T	3, 0	1, 2
B	0, 4	0, -100

Klibanoff shows that (T, L) is a KE. To see this let $C_1 = [0.1, 1]$ and $C_2 = \{1\}$. Player 2's utility of his strategy p_2 can be written as $\min_{q \in [0.1, 1]} 2(1 - p_2)q + 4p_2(1 - q) - 100(1 - p_2)(1 - q)$. Player 2's utility is maximized at $p_2 = 1$. On the other hand T is player 1's best reply to the pure strategy L . Next, observe that (T, L) and C_1, C_2 satisfy the consistency requirement. Finally, observe that the unique NE of the game is (T, R) .

To understand better the difference between the concepts of AAE and KE, let me rewrite the above example to fit the present framework: consider a general ambiguous act extension $G' = (\Omega, A, \succ')$ of the above game G with $\Omega_1 = \{\omega_1, \omega_2\}$ and Ω_2 a singleton. Define f such that $f_1(\omega_1) = T$, $f_1(\omega_2) = B$ and $f_2 = L$. For this profile to be an equilibrium ω_2 has to be 1-null as T strictly dominates B . On the other hand, ω_2 cannot be 2-null, as in that case R would be player 2's unique best reply. We conclude that players must disagree on null events for (T, L) to be an equilibrium. In accordance with the section on observational equivalence (Section 6), the connection between this strategy profile (T, L) to some empirical predictions becomes difficult: if the researcher was to agree with player 1 he should be utterly surprised if B was ever observed. If, on the other hand, he agrees more with player 2's view of the world, the play of B should be perceived as normal.

Next, it is important to observe that the assumption that players disagree on null events alone can generate the same strategy profile – uncertainty aversion is not needed. The acts f_1, f_2 are best replies to each other when both players are expected utility maximizers with divergent priors $\pi_1(\omega_1) = 1$ and $\pi_2(\omega_1) = 1/10$. Finally, it should be noted that (MON) needs to be violated for (T, L) to be a KE, as the consistency condition requires that $p_1 = 1 \in C_1$ which must contain other priors that assign positive probability to ω_2 as was argued above.

The equilibrium concepts of Lo and Klibanoff differ only insofar as Lo requires a stronger consistency condition between strategies and beliefs. This condition is strong enough to imply that players must agree on non-null events; Lo therefore also obtains a result of observational equivalence. Interestingly, disagreement on null events also underlies most examples provided by Dow and Werlang and Marinacci, even though their concepts could generate observationally different results without such disagreement as they extend the notion of equilibrium in beliefs to the case of uncertainty-averse players (see the discussion in Section 11.3).

12. Conclusion

The first contribution of this paper is to put Aumann's (1974) framework of games into service for the analysis of strategic interactions between ambiguity-averse players. This framework allows us to base the definition of an ambiguous analogue to mixed-strategy equilibrium on the standard notion of Nash equilibrium. The second main contribution of the paper is the proof that an outside observer cannot distinguish whether a game is played by two uncertainty-averse players or two uncertainty-neutral ones if players use explicit randomization devices. The third main contribution concerns the interpretation of alternative equilibrium concepts for uncertainty-averse players. I show that the behavioral differences predicted by prior equilibrium concepts in the literature are not so much a result of the assumption of uncertainty aversion, but rather a result on the players disagreements on the possible occurrence of all events in the state space. Is there any

hope for a manageable theory of games with uncertainty-averse players that yields predictions that differ from standard theory?

For me, the answer is a clear yes. I see the following avenues for future research. We might want to give up on the assumption that players should agree on null events. If this is the case, then we should do so in a controlled manner: Marinacci (2000) does exactly that. The advantage of his equilibrium concept is that he parameterizes the uncertainty of players in a game. In the light of the present study, such a parametrization seems very important as it allows us gradually to relax the assumption that players agree on null events. Marinacci's approach allows us to find equilibrium predictions for ambiguity-averse players that differ from the equilibrium predictions of mixed-strategy equilibrium while retaining control over the gap between the players' actual strategies and other players' beliefs on these strategies.⁸ Marinacci's main contribution is a proof of existence of ME for any level of uncertainty. The concept has yet to prove its merits in applied studies.

Even if we insist on players agreeing on null events, a game theory with uncertainty-averse players might yield observationally different results. Theorem 1 crucially depends on the assumption that there are just two players. In Section 11.2, I showed that the extension of the present theory to more than two players can yield observationally different results. Schmeidler's (1989) assumption of uncertainty aversion (UA) also plays a certain role for the validity of Theorem 1. If games with ambiguity-averse players are instead modeled following Bewley (2002), predictions that differ observationally from the standard theory of mixed-strategy equilibrium can be obtained. In Section 11.3, I showed that a translation of the concept of equilibrium in beliefs to the context of ambiguity-averse players can yield observationally different results from the concept of mixed-strategy equilibrium. In short: the negative result for the model of two-player games introduced here should highlight the potential for all the positive results that could arise when relaxing the various assumptions made in the present model.

References

- Aumann, R., 1974. Subjectivity and correlation in randomized strategies. *J. Math. Econ.* 1, 67–96.
- Azrieli, Y., Teper, R., 2009. Uncertainty aversion and equilibrium existence in games with incomplete information. Mimeo, Tel Aviv University.
- Bade, S., 2007. Electoral competition with uncertainty averse parties. Mimeo, Penn State University.
- Bade, S., 2008. Stochastic independence with maxmin expected utilities. Mimeo, Penn State University.
- Bewley, T., 2002. Knightian decision theory: Part I. *Dec. Econ. Finance* 25, 79–110.
- Billot, A., Chateauneuf, A., Gilboa, I., Tallon, J.-M., 2000. Sharing beliefs: Between agreeing and disagreeing. *Econometrica* 68, 685–694.
- Bose, S., Ozdenoren, E., Pape, A., 2006. Optimal auctions with ambiguity. *Theoretical Econ.* 1, 411–438.
- Brandenburger, A., Blume, L., Dekel, E., 1991. Lexicographic probabilities and choice under uncertainty. *Econometrica* 59, 61–79.
- Camerer, C., Weber, M., 1992. Recent developments in modeling preferences: Uncertainty and ambiguity. *J. Risk Uncertainty* 5, 325–370.
- Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L., 2008. Uncertainty averse preferences. *Carlo Alberto Notebook* 77.
- Chiappori, P.-A., Levitt, S., Grogg, T., 2002. Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *Amer. Econ. Rev.* 92, 1138–1151.
- Dow, J., Werlang, S., 1994. Nash equilibrium under knightian uncertainty: Breaking down backward induction. *J. Econ. Theory* 64, 305–324.
- Eichberger, J., Kelsey, D., 2000. Non-additive beliefs and strategic equilibria. *Games Econ. Behav.* 30, 183–215.
- Epstein, L., 1997. Preference, rationalizability and equilibrium. *J. Econ. Theory* 73, 1–29.
- Epstein, L., Marinacci, M., 2007. Mutual absolute continuity of multiple priors. *J. Econ. Theory* 137, 716–720.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *J. Math. Econ.* 18, 141–153.
- Hansen, L., Sargent, T., 2001. Robust control and model uncertainty. *Amer. Econ. Rev.* 91, 60–66.
- Klibanoff, P., 1996. Uncertainty, decision and normal form games. Mimeo, Northwestern.
- Klibanoff, P., 2001. Stochastically independent randomization and uncertainty aversion. *Econ. Theory* 18, 605–620.
- Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Lehrer, E., 2008. Partially-specified probabilities: Decisions and games. Mimeo, Tel Aviv.
- Levin, D., Ozdenoren, E., 2004. Auctions with uncertain numbers of bidders. *J. Econ. Theory* 118, 229–251.
- Lo, K.C., 1996. Equilibrium in beliefs under uncertainty. *J. Econ. Theory* 71, 443–484.
- Maccheroni, F., Marinacci, M., Rustichini, A., 2006. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74, 1447–1498.
- Marinacci, M., 2000. Ambiguous games. *Games Econ. Behav.* 31, 191–219.
- Mukerji, S., Tallon, J.-M., 2004. An overview of economic applications of David Schmeidler's models of decision making under uncertainty. In: Gilboa, I. (Ed.), *Uncertainty in Economic Theory: A Collection of Essays in Honor of David Schmeidler's 65th Birthday*. Routledge Publishers (Chapter 13).
- Palacios-Huerta, I., 2003. Professionals play minimax. *Rev. Econ. Stud.* 70, 395–415.
- Schmeidler, D., 1989. Subjective probability and expected utility without additivity. *Econometrica* 57, 571–587.
- Walker, M., Wooders, J., 2001. Minimax play at Wimbledon. *Amer. Econ. Rev.* 91, 1521–1538.

⁸ Eichberger and Kelsey (2000) provide an alternative parametrization of the degree of uncertainty in a game.