TITLE:

Savage Games

A Theory of Strategic Interaction with Purely Subjective Uncertainty

Authors:

Simon Grant, Idione Meneghel and Rabee Tourky

Working Paper: R13_3

2013

RSMG Working Paper Series

Risk and Uncertainty Program

Schools of Economics and Political Science

The University of Queensland

St Lucia

Brisbane

Australia 4072

Web: www.uq.edu.au

SAVAGE GAMES
A Theory of Strategic Interaction with Purely Subjective Uncertainty

SIMON GRANT, IDIONE MENEGHEL, AND RABEE TOURKY

Abstract. We define and discuss Savage games, which are ordinal games that are set in L. J. Savage's framework of purely subjective uncertainty. Every Bayesian game is ordinally equivalent to a Savage game. However, Savage games are free of priors, probabilities and payoffs. Players' information and subjective attitudes toward uncertainty are encoded in the state-dependent preferences over state contingent action profiles. In the games we study player preferences satisfy versions of Savage's *sure thing principle* and *small event continuity* postulate. An axiomatic innovation is a strategic analog of Savage's null events. We prove the existence of equilibrium in Savage games. This result eschews any notion of objective randomization, convexity, and monotonicity. Applying it to games with payoffs we show that our assumptions are satisfied by a wide range of decision-theoretic models. In this regard, Savage games afford a tractable framework to study attitudes towards uncertainty in a strategic setting. We illustrate our results on the existence of equilibrium by means of examples of games in which players have expected and non-expected utility.

Keywords: Bayesian games, multiple priors, non-expected utility, subjective uncertainty, existence of equilibrium, decomposable sets.

1. Introduction

Consider the N-player Bayesian game

$$((\Omega_i, \Sigma_i), A_i, u_i, \pi)_{i=1}^N ,$$

where (Ω_i, Σ_i) is the measurable space of Player i's *types* and A_i is a compact metric space of Player i's *actions*. Player i has a bounded measurable *payoff function*

$$u_i \colon A \times \Omega \to \mathbb{R} ,$$

where $A = \times_j A_j$ is the set of *action profiles* and $\Omega = \times_j \Omega_j$ is the space of *type profiles*. Letting F_j be the set of all measurable *strategies* $f_j \colon \Omega_j \to A_j$ of Player j, Player i's expected utility for the *strategy profile* f in $F = \times_j F_j$ is

$$U_i(f) = \int_\Omega u_i(f(\omega), \omega) \, \mathrm{d}\, \pi(\omega) ,$$

where π is the probability measure on type profiles representing the players' common *prior*.

This induces for Player i a preference relation $\succsim_i$ on the set of strategy profiles F, whereby $f \succsim_i g$ whenever $U_i(f) \geq U_i(g)$. The triple, $(\Omega, A, \succsim_i)$, can in turn be articulated in terms of Savage's (1954) framework of purely subjective uncertainty. The set of type profiles Ω is Savage's state space; the set of action profiles A plays the role of Savage's outcome space; and preferences $\succsim_i$ are restricted to those Savage acts $f \colon \Omega \to A$ that are strategy profiles.

Our interest is to study N-player games specified by $(\Omega, A, (F_i, \succsim_i)_{i=1}^N)$, where Ω is the common state space and A is the common action space. For Player i, F_i is a given subset of A-valued functions on Ω comprising those strategies available to that player and $\succsim_i$ is

Date: June 18, 2013.

We are grateful for helpful comments from Rabah Amir, Jean-Marc Bonnisseau, Larry Epstein, Michael Greinecker, Atsushi Kajii, George Mailath, Roee Teper and Nicholas Yannelis.

a preference relation defined over strategy profiles in $F = \times_j F_j$. We call these *Savage games*. They redact all cardinal aspects of Bayesian games (utilities and priors), the intrinsic attitudes towards uncertainty are purely subjective, and the common state and action spaces need not be Cartesian products of underlying individual state and action spaces. An *equilibrium* is a Nash equilibrium of the N-player ordinal game $(F_i, \succsim_i)_{i=1}^N$.

We study three behavioral properties that are inherited from Bayesian games.

(a) Any deviation by a player from one best response to another conditional on an event remains a best response. This version of Savage's *sure thing principle* holds in all Bayesian games.
(b) If a deviation by Player i affects Player j's preferences, then the effect on j is robust to Player i making "errors" that are conditioned on the *small events* specified by some partition of the state space. This interdependent version of Savage's *small event continuity* postulate is satisfied in a continuous Bayesian game with a common prior that has atomless marginals.
(c) Analogously to Savage's *null* events, we say that an event is *strategically null* for Player i if any deviation on that event by the player does not affect any players' preferences. Events that are not null are "full" in a behavioral sense: for every sequence of events that has no subsequence decreasing to a null event there is a state whose average incidence is asymptotically bounded away from zero. In a Bayesian game with common priors every zero probability event is a strategically null event; hence, this property is satisfied.

With the following additional condition we prove the existence of equilibrium and in the process develop a theory of equilibrium existence that is derived entirely from the ordinal preferences of the players.

(d) Players have the opportunity to respond to other players using strategies that are in an equicontinuous class of functions. In a Bayesian game, this is satisfied when one can always find a best response with bounded marginal change in chosen actions with respect to own types.

Turning to games with payoffs, we study properties (a), (b), and (c) when the *ex ante* evaluation of strategy profiles can be expressed as a recursive function of interim utilities. We show that a wide range of properties studied in the literature satisfy our assumptions. The analysis includes games in which players' *ex ante* utilities are of the *maxmin* expected utility form of Gilboa and Schmeidler (1989) if the following three conditions hold.

(1) Payoffs are continuous in action profiles.
(2) There is a measure with atomless marginals that dominates the multiple priors of all players.
(3) The preferences admit a utility representation in the recursive form of Epstein and Schneider (2003) and for each player the multiple priors are mutually absolutely continuous (cf. Epstein and Marinacci (2007)).

Here an application of our existence theorem simply requires us to check condition (d). We illustrate this by establishing the existence of equilibrium in two versions of a location model on the n-sphere. The first is a Bayesian game in which each player has her own (not necessarily common) prior and the second is a game with multiple priors for each player. In these examples players simultaneously choose locations on an n-dimensional unit sphere and have payoff functions that depend in a general way on the location choices and types of all players. We establish that property (d) holds by constructing a parametrization of strategies using a selection theorem for closed-valued Lipschitz continuous correspondences.

The theorem on the existence of equilibrium develops techniques and utilizes a new fixed point theorem in non-linear analysis (for this, see the discussion following Theorem 2.8 and Theorem A.4.) Our techniques extend the concepts in Athey (2001), McAdams (2003) and Reny (2011), who proved the existence of monotone equilibrium in Bayesian games with continuous lattice-ordered action spaces and atomless type spaces. The recent paper of Reny (2011) provides the most refined argument, he uses a homotopy method to trace decomposition paths in the sub-semilattices of monotone best responses. The present work frees these techniques from any order-theoretic requirement and extends the approach well beyond Bayesian games. The key idea is that the set of strategies and the sets of best responses are decomposable in the sense of Rockafellar (1968, 1971) and related literature on non-linear analysis. This decomposable choice property was studied in a decision-making framework by Grant, Kajii, and Polak (2000), who applied their own tracing arguments. By means of conditional deviations, in this paper we can continuously pass from one best response to another without leaving the set of best responses. This property, which is a consequence of our *sure thing principle*, our *small event continuity*, and the "fullness" property of non-null sets (respectively, (a), (b), (c) above) allows us to apply path following arguments without any convexity assumption, or any meaningful notion of monotonicity. Assumption (d) ensures that we have the required compactness property to establish the existence of equilibrium.

The paper is organized as follows. In Section 2 we describe and study Savage games and state our main theorem. In Section 3 we study games with recursive payoffs and priors. We highlight in this section how our assumptions and the result on the existence of equilibrium translate to Bayesian games, games with multiple priors and games in which preferences display other forms of non-expected utility. Section 4 focuses specifically on preferences that have the *maxmin* expected utility representation. Section 5 contains the two examples. We conclude with a discussion of open questions and possible extensions of this work in Section 6. The proofs are in Appendix A.

2. Savage games

A *Savage game* is an N-player ordinal game modeling choice under uncertainty with interdependent preferences. It is specified by the ternion

$$(\Omega, A, (F_i, \succsim_i)_{i=1}^N) .$$

The set Ω denotes the common *state space* and A is the common nonempty *action* set endowed with a topology in which single points are closed. The tuple $(F_i, \succsim_i)_{i=1}^N$ is an N-player ordinal game whose parameters are described below.

There are $N \geq 1$ players indexed by $i = 1, \dots, N$. We abuse notation by having N also denote the set $\{1, \dots, N\}$. However, we employ standard notation for the indexing of player profiles. In particular, for any N-tuple $(Z_i)_{i=1}^N$ of sets we write Z for its N-ary Cartesian product and for each player i we write Z_{-i} for the Cartesian product of the tuple $(Z_j)_{j\neq i}$. Vectors in Z are called *profiles* and vectors in Z_{-i} are called *profiles of players other than* i. A profile $z \in Z$ is also written as (z_i, z_{-i}) where z_i is the i-th coordinate of z and z_{-i} is the projection of z into Z_{-i}.

Player i has a non-empty set F_i of A-valued functions on the state space Ω called the *strategy space*. A function $f_i \colon \Omega \to A$ in F_i is called a *strategy* for player i. Let F be the set of strategy profiles and for each i let F_{-i} be the set of strategy profiles of players other than i.

Player i is also associated with a preordering $\succsim_i$ on the set of strategy profiles F describing her weak preferences. That is, $\succsim_i$ is transitive and reflexive. Let $\sim_i$ represent the indifference relation associated with $\succsim_i$, that is $f \sim_i g$ if $f \succsim_i g$ and $g \succsim_i f$. We assume throughout that the following completeness condition is satisfied.

A1. For any $f \in F$, $g_i \in F_i$ either $f \succsim_i (g_i, f_{-i})$ or $(g_i, f_{-i}) \succsim_i f$.

A strategy profile $f \in F$ is an *equilibrium* if

$$f \succsim_i (g_i, f_{-i})$$

for all $g_i \in F_i$ and $i \in N$. Under A1 $f \in F$ is an equilibrium if and only if it is a Nash equilibrium of the ordinal game $(F_i, \succsim_i)_{i=1}^N$.

In a Savage game the information available to a player is encoded in the specification of the set of strategies F_i. Following standard notation, for any subset $E \subseteq \Omega$ and two functions $f_i, g_i \colon \Omega \to A$ let $g_{iE}f_i$ be the function from Ω to A given by

$$g_{iE}f_i(\omega) = \begin{cases} g_i(\omega) & \text{if } \omega \in E\,, \\ f_i(\omega) & \text{otherwise}\,. \end{cases}$$

We refer to the function $g_{iE}f_i$ as the g_i-deviation from f_i conditional on E.

Information events. A set of states $E \subseteq \Omega$ is an (information) *event* for Player i if she can condition her choice of strategy on E, that is, $g_{iE}f_i \in F_i$ for all $f_i, g_i \in F_i$. Denote by $\mathcal{F}_i$ the family of events for Player i.

The next assumption is motivated by Savage's *sure thing principle.* We shall omit the quantifiers from our assumptions when they are obvious. In particular, f is understood as an arbitrary member of F, f_i and g_i of F_i, and E always denotes an event in $\mathcal{F}_i$.

A2. If $f \sim_i (g_i, f_{-i}) \succsim_i (g_{iE}f_i, f_{-i})$ for all $E \in \mathcal{F}_i$, then $(g_{iE}f_i, f_{-i}) \sim_i f$ for all $E \in \mathcal{F}_i$.

The reader will see in Section 3 that in games in which the preferences are given by payoffs A2 holds when any of a wide range of assumptions that have been studied in the literature are satisfied. In particular, A2 is implied by Savage's postulate **P2**.

Proposition 2.1. The following condition implies A2

P2: If $f \succsim_i (g_{iE}f_i, f_{-i})$, then $(f_{iE}g_i, f_{-i}) \succsim_i (g_i, f_{-i})$.

We do not assume that the game contains constant strategies or that it is non-degenerate in the sense of Savage. We make, however, the following "richness" assumption on strategies.

A3. If $E^n \in \mathcal{F}_i$ is an increasing sequence, then $g_{i(\cup_n E^n)}f_i \in F_i$ for all $f_i, g_i \in F_i$.

Assumption A3 guarantees that $\mathcal{F}_i$ is a σ-algebra.

Proposition 2.2. $\mathcal{F}_i$ is an algebra over Ω. If A3 holds, then it is a σ-algebra.

When a measure-theoretic framework is available we have the following.

Corollary 2.3. Let Σ_i be a σ-algebra over Ω and A be a compact metric space, with $|A| \geq 2$. If F_i is the set of all Σ_i-measurable functions to A, then A3 holds and $\mathcal{F}_i = \Sigma_i$.

We extend the concept of a Savage null event to our setting of interdependent preferences. An event will be deemed (Savage) null for a player if any deviation that player can make conditional on that event from any strategy profile leaves *all* players indifferent.

Null events. An event $E \in \mathcal{F}_i$ is *null* for Player i if for all $f \in F$ and all $g_i \in F_i$ we have $(g_{iE}f_i, f_{-i}) \sim_j f$, for every player $j \in N$. Denote by $\mathcal{N}_i$ the set of all events that are null for Player i. Let $\mathcal{R}_i = \mathcal{F}_i \setminus \mathcal{N}_i$ be the set of *relevant* events for Player i.

Notice that two players $i, j \in N$ may share an event $E \in \mathcal{F}_i \cap \mathcal{F}_j$ that is null for Player i but relevant for Player j. We do not view this as anomalous or inconsistent. It simply means that when conditioning on this event, no deviation by Player i can make anyone better or worse off, however, there exists at least one deviation by Player j that makes at least one player in the game either better or worse off.[1]

[1] Karni, Schmeidler, and Vind (1983) note that if preferences are state-dependent, then interpreting null events as ones that are necessarily viewed by the decision-maker as having zero probability of occurring is problematic. For example, if one of the events involves loss of life then its nullity could reflect the decision-maker having no strict preference about which outcome obtains in the event she is dead, rather than her believing she has no chance of dying.

We turn now to the continuity of preferences. While the product topology is too coarse to be useful for our purpose, a fine topology that is closely related to metric spaces is the *sequential topology*, whereby a set W of functions from Ω to A is open if every sequence of functions converging state-wise to a function in W is eventually in W.[2] We require continuity of preferences in the sequential topology, and note that the condition is satisfied in continuous games in which utility is computed by means of integrals.

A4. If $f^n \in F$ and $g_i^n \in F_i$ are two sequences converging state-wise to $f \in F$ and $g_i \in F_i$, respectively, and $(g_i^n, f_{-i}^n) \succsim_i f^n$ for all n, then $(g_i, f_{-i}) \succsim_i f$.

The set of null events for Player i is automatically an ideal, and assumption A4 guarantees that it is a σ-ideal.

Proposition 2.4. $\mathcal{N}_i$ is an ideal in $\mathcal{F}_i$. If A3 and A4 hold for all players, then $\mathcal{N}_i$ is a σ-ideal.

The next assumption is a "fullness" assumption on relevant events. It adapts the Ryll-Nardzewski and Kelley condition for Boolean algebras to our setting (see the addendum to Kelley (1959)). A family of events $\mathcal{S} \subseteq \mathcal{F}_i$ is *closed* if for any increasing sequence of events E^n in $\mathcal{S}$ whose union E is an event, we have $E \in \mathcal{S}$.

A5. There is a sequence of closed families $\mathcal{S}_i^m$ of events satisfying:

(1) If $E \in \cap_m \mathcal{S}_i^m$, then E is null for Player i.

(2) If E^n is a sequence of relevant events for Player i and

$$\liminf_{n\to\infty} \quad \max_{\omega\in\Omega} \tfrac{1}{n}|\{1 \le k \le n \colon \omega \in E^k\} = 0\,,$$

then for each m there is n such that $E^n \in \mathcal{S}_i^m$.

Interpreting the sequence $\mathcal{S}_i^m$ as families of small events forming a neighborhood base for the subfamily of null sets in condition (1), condition (2) asserts that if the average incidences arising from the sequence E^n uniformly converges to zero for all states, then E^n has a subsequence that "converges" to a null event. Condition (2) guarantees that for each $\mathcal{S}_i^m$ there is a number $c_m > 0$ such that if X^n is a sequence of random variables and for some $\alpha \in \mathbb{R}$ the sets $\{\omega \colon X^n(\omega) \ge \alpha\}$ are not in $\mathcal{S}_i^m$, then the empirical cumulative distribution of this sequence satisfies $D_X(\beta) \le 1 - c_m$ for every $\beta < \alpha$.

The following result is proved in Appendix A.

Proposition 2.5. If A3, A4 hold for all players, then the following are equivalent:

(1) Assumption A5 holds for Player i.

(2) $\mathcal{F}_i$ admits a measure π_i such that $\pi_i(E) = 0$ if and only if $E \in \mathcal{N}_i$.

The next assumption is an interdependent version of Savage's postulate **P6**, which is usually interpreted as a *small event continuity* property.

A6. If $f \not\sim_j (g_i, f_{-i})$ for some $j \in N$, then for each $h_i \in F_i$ there exist events $\{E^1, \dots, E^k\}$ such that $\cup_k E^k = \Omega$ and $f \not\sim_j (h_{iE_n} g_i, f_{-i})$ for all k.

As with **P6**, this assumption ensures that relevant events are the union of two disjoint relevant events.

Proposition 2.6. If A6 holds, then every $E \in \mathcal{R}_i$ contains two disjoint events in $\mathcal{R}_i$.

We next characterize when A5 and A6 hold. The proof of the next proposition is also in Appendix A.

[2] With small event continuity, which appears in the sequel, the continuity of preferences in the product topology, involving convergent nets of strategies, implies that $f \sim_i g$ for all pairs of strategy profiles. Notice that g_i is the limit of the net $g_{iE} f_i$ where $E \in \mathcal{N}_i$, directed by inclusion.

Proposition 2.7. If A3 and A4 hold for all players, then the following are equivalent:

(1) Assumptions A5 and A6 hold for Player i.
(2) $\mathcal{F}_i$ admits an atomless measure π_i such that $\pi_i(E) = 0$ implies $E \in \mathcal{N}_i$.

Notice that in condition (2) some null events of a player may be given a positive measure. The only requirement is that all events assigned zero measure are null for that player. Since we are dealing with countably additive measures, however, there exists a corresponding measure that is positive on relevant events and zero on all null events.

State-wise agreeing strategies. A strategy $f_i \in F_i$ *agrees* with a set of strategies $X_i \subseteq F_i$ if

$$f_i = g^1_{iE^1}(g^2_{iE^2}(\cdots(g^{K-1}_{iE^{K-1}}g^K)))$$

for some $g^k_i \in X_i$ and $E^k \in \mathcal{F}_i$, $k = 1, \ldots, K$. Denote by $\widetilde{X}_i$ the closure of $\{f_i \in F_i \colon f_i \text{ agrees with } X_i\}$ in the sequential topology. Functions in $\widetilde{X}_i$ are said to *state-wise agree* with X_i.

A subset of strategies $X_i \subseteq F_i$ is said to be *countably distinguished* if there exists a countable set of states $\mathbb{W} \subseteq \Omega$ such that for any distinct $f_i, g_i \in X_i$ we have $f_i(\omega) \neq g_i(\omega)$ for some $\omega \in \mathbb{W}$.

A7. For each i there is a set $X_i \subseteq F_i$ of strategies satisfying the following:

(1) For each $f_i \in F$, each $f_{-i} \in \widetilde{X}_{-i}$ there is $g_i \in X_i$ satisfying $(g_i, f_{-i}) \succsim_i f$.
(2) X_i is countably-distinguished and every sequence in X_i has a subsequence converging state-wise to a strategy in X_i.

Assumption A7(2) is equivalent to saying that X_i is metrizable and compact in the sequential topology. For clarity, we provide some examples of such spaces.

(a) X_i is a finite set.
(b) X_i contains a countable set and its countable accumulation points.
(c) If Ω and A are compact metric spaces, then any closed collection of equicontinuous strategies satisfies this assumption. In particular, X_i is the set of all Lipschitz continuous functions with common constant.
(d) Suppose that Ω is the set of all continuous functions from $[0, 1]$ to $\mathbb{R}$, and A is the set of Radon measures on $[0, 1]$. Then the set of all Radon probability measures with the weak*-topology has the required property.
(e) Suppose that $\Omega = [0, 1]$ and $A = \mathbb{R}^d$. For any compact set of functions Y in L_∞, there is a selection from the equivalence classes of these functions that is compact and metrizable in the sequential topology.
(f) Suppose that Ω is a measure space with σ-algebra Σ, and A is a Banach space. This assumption is satisfied whenever X_i is a sequentially compact set of bounded measurable functions for the topology of state-wise convergence, and there is a probability measure π on Σ such that if f_i and g_i are distinct functions in X_i, then they differ on a set of positive π measure.

We are now ready to state the main result.

Theorem 2.8. If A1 to A6 hold for all players, then an equilibrium exists if and only if A7 holds.

The proof is in Appendix A but we conclude this section with a brief overview of the line of argument. The proof relies on a fixed point theorem for closed-graphed decomposable-valued correspondences. The idea is a novel notion of monotone purification that allows us to apply a fixed point theorem for absolute retracts. Basically, we take the subset of strategy profiles X given by assumption A7. While X does not admit any continuous ordering, it is compact and metrizable, and thus it is the image of a closed subset C of $[0, 1]$. We then consider the set of cumulative distributions on C, and show that each of these cumulative distributions maps back to a strategy profile in $\widetilde{X}$. So we transfer the existence of equilibrium problem to a fixed point problem in the space of cumulative distributions on C. By means of A1 and A3 to A6, we transform the best response correspondence into a closed-graphed correspondence on our set of cumulative distributions. A2 guarantees two things. First, that any fixed point of the correspondence on cumulative distributions can be mapped back by a "monotone purification argument" into an equilibrium of the game. Second, it guarantees that the transformed best response correspondence is an acyclic correspondence on an absolute retract, which has a fixed point by the Eilenberg and Montgomery theorem.

3. Games with recursive payoffs

Our purpose in this section is to study a class of games in which a player's *ex ante* evaluation of strategy profiles can be expressed as a recursive function of her *interim* utilities or payoffs. The assumptions that we impose on these games essentially allow for almost arbitrary assessments by a player of other players' strategies. The Savage games associated with these games naturally satisfy assumptions A1 to A6. The games include Bayesian games, games with multiple priors that admit a recursive representation, and more generally games with preferences over interim payoffs satisfying many conditions studied in the literature on decision making under uncertainty. We shall study the specific class of games with multiple priors in the next section.

Consider a game in *interim utility* form specified as follows:

$$((\Omega_i, \Sigma_i), A_i, \mathbf{V}_i, \mathbf{U}_i)_{i=1}^N ,$$

where (Ω_i, Σ_i) is the measurable space of Player i's *types* and A_i is a compact metric space of Player i's *actions.* The space of *type profiles* $\Omega = \times_{i=1}^N \Omega_i$ has the product algebra $\Sigma = \otimes_{i=1}^N \Sigma_i$.

Let F_i be the set of all Σ_i-measurable *strategies* $f_i \colon \Omega_i \to A_i$. Player i is associated with an *interim utility function* $\mathbf{V}_i \colon A_i \times F_{-i} \times \Omega_i \to \mathbb{R}$, where $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ is the interim utility for Player i whose type is ω_i if she chooses action a_i when the other players are choosing their actions according to the strategy profile f_{-i}. For any strategy profile $f \in F$ we write $\mathbf{V}_i(f)$ for the real-valued function $\omega_i \mapsto \mathbf{V}_i(f_i(\omega_i), f_{-i}|\omega_i)$, which we assume is always bounded and Σ_i-measurable.

We call a bounded and Σ_i-measurable real-valued function $\alpha_i \colon \Omega_i \to \mathbb{R}$ an *interim payoff* for Player i. Player i has *ex ante* preferences over interim payoffs expressed by the utility function $\mathbf{U}_i$ that associates with each α_i an *ex ante* utility $\mathbf{U}_i(\alpha_i)$ in $\mathbb{R}$. The *ex ante* utility $U_i(f)$ of Player i for the strategy profile $f \in F$ is given by means of the recursive form

$$U_i(f) = \mathbf{U}_i \circ \mathbf{V}_i(f) .$$

An *equilibrium* is a Nash equilibrium of the normal form game $(F_i, U_i)_{i=1}^N$.

Extending the notation of the previous section to interim payoffs we make the following decomposition assumption on preferences over interim payoffs.

B1. If α_i, β_i are two interim payoffs and $\mathbf{U}_i(\alpha_i) = \mathbf{U}_i(\beta_i) \geq \mathbf{U}_i(\beta_{iE_i}\alpha_i)$ for all $E_i \in \Sigma_i$, then $\mathbf{U}_i(\beta_{iE_i}\alpha_i) = \mathbf{U}_i(\alpha_i)$ for all $E_i \in \Sigma_i$.

We also require that *ex ante* utility over strategy profiles be continuous with respect to an atomless measure.

B2. There exists a probability distribution $\mu \colon \Sigma \to [0, 1]$ such that

(1) All the marginal distributions $\hat{\mu}_i \colon \Sigma_i \to [0, 1]$ of μ are atomless.
(2) If f^n is a sequence of strategy profiles that converges μ-almost everywhere to f, then $U_i(f^n)$ converges to $U_i(f)$ for all i.

Let $\bar{A}$ be the disjoint union of the sets A_i endowed with a consistent metric for which it is compact. For each $f_i \in F_i$ let $\bar{f}_i$ be the $\bar{A}$-valued function on Ω given by $\bar{f}_i(\omega) = f_i(\omega_i)$. Define $\bar{F}_i = \{\bar{f}_i \colon f_i \in F_i\}$, which yields the Savage game

$$(\Omega, \bar{A}, (\bar{F}_i, \succsim_i)_{i=1}^N)$$

where $\succsim_i$ is given by $\bar{f} \succsim_i \bar{g}$ if and only if $U_i(f) \geq U_i(g)$ for any $\bar{f}, \bar{g} \in \bar{F}$.

Proposition 3.1. *If B1 and B2 hold, then the associated Savage game satisfies A1 to A6.*

We now explore properties of players' *ex ante* attitudes towards interim payoffs as embodied in $\mathbf{U}_i$, which guarantee that B1 is satisfied. Of course, $\mathbf{U}_i$ only depends on a Player i's own type, so behaviorally the properties that we discuss are purely decision theoretic embodying the player's attitudes toward non-strategic uncertainty. In this regard, these properties can be compared to the generalizations of expected utility in the literature.

Example 3-1 (Strictly monotone utility)
The first property simply requires that players have strictly monotonic preferences over interim payoffs. Some form of monotonicity is present in nearly all generalizations of expected utility. Using the measure from B2, if $\mathbf{U}_i$ is strictly monotone for the marginal $\hat{\mu}_i$-pointwise ordering of interim payoffs, then B1 holds.

Let μ be the measure from B2. For any interim payoffs write $\alpha_i \geq \beta_i$ if $\alpha_i(\omega_i) \geq \beta_i(\omega_i)$ for $\hat{\mu}_i$-almost all ω_i. Write $\alpha_i > \beta_i$ if $\alpha_i \geq \beta_i$ and $\alpha_i(\omega_i) > \beta_i(\omega_i)$ over a set of positive $\hat{\mu}_i$-measure. Suppose that $\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\beta_i)$ holds whenever $\alpha_i \geq \beta_i$ and $\mathbf{U}_i(\alpha_i) > \mathbf{U}_i(\beta_i)$ whenever $\alpha_i > \beta_i$. We show that B1 holds. Take α_i, β_i from that condition. Notice that

$$\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\alpha_i \vee \beta_i) \geq \mathbf{U}_i(\alpha_i)\,,$$

where $\alpha_i \vee \beta_i$ is the pointwise supremum of the interim payoffs. Thus, $\alpha_i \vee \beta_i(\omega_i) = \alpha_i(\omega_i)$ $\hat{\mu}_i$-almost surely. This implies that $\alpha_i(\omega_i) = \beta_{iE_i}\alpha_i(\omega_i)$ $\hat{\mu}_i$-almost surely for any $E_i \in \Sigma_i$. Monotonicity now tells us that $\mathbf{U}_i(\alpha_i) = \mathbf{U}_i(\beta_{iE_i}\alpha_i)$, as required. ■

Example 3-2 (Supermodular utilities)
In the presence of ambiguity aversion, preferences over interim payoffs need not be strictly monotonic though weak monotonicity can usually be guaranteed (see for example Gilboa (1987)). However, if $\mathbf{U}_i$ can be represented by a Choquet integral and exhibits Schmeidler's (1989) notion of uncertainty aversion, then $\mathbf{U}_i$ will be supermodular (Denneberg, 1994, Corollary 13.4, p. 161). In fact, weak monotonicity and supermodularity together imply that condition B1 holds.

For any interim payoffs α_i and β_i denote by $\alpha_i \vee \beta_i$ and $\alpha_i \wedge \beta_i$ the state-wise supremum and infimum payoffs. Assume that $\mathbf{U}_i$ is non-decreasing in the sense that if $\alpha_i(\omega_i) \geq \beta_i(\omega_i)$ for all ω_i, then $\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\beta_i)$. Now suppose that $\mathbf{U}_i$ satisfies supermodularity:

$$\mathbf{U}_i(\alpha_i \vee \beta_i) + \mathbf{U}_i(\alpha_i \wedge \beta_i) \geq \mathbf{U}_i(\alpha_i) + \mathbf{U}_i(\beta_i)\,.$$

We show that B1 holds. Take α_i, β_i from that condition and notice that

$$\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\alpha_i \vee \beta_i) \text{ and } \mathbf{U}_i(\beta_i) \geq \mathbf{U}_i(\alpha_i \wedge \beta_i)\,,$$

because $\alpha_i \vee \beta_i = \beta_{iE}\alpha_i$ and $\alpha_i \wedge \beta_i = \alpha_{iE_i}\beta_i$ for $E = \{\omega_i \colon \beta_i(\omega_i > \alpha_i)\}$. Therefore, $\mathbf{U}_i(\alpha_i \wedge \beta_i) = \mathbf{U}_i(\alpha_i)$. From this and the monotinicity of $\mathbf{U}_i$ we conclude that for any $E \in \Sigma_i$ we have

$$\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\beta_{iE}\alpha_i) \geq \mathbf{U}_i(\alpha_i \wedge \beta_i) = \mathbf{U}_i(\alpha_i)\,,$$

as required. ■

Example 3-3 (Decomposable choice)
Moving away from explicit monotonicity, we give a "betweenness" condition on preferences over interim payoffs that generalizes Savage's P2 postulate and that also satisfies the property B1 above. Consider the *decomposable choice* property of Grant, Kajii, and Polak (2000), which in this setting requires the following:

GKP: For any interim payoffs α_i, β_i and events $E \in \Sigma_i$ if $\mathbf{U}_i(\alpha_i) > \mathbf{U}_i(\beta_{iE}\alpha_i)$ and $\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\alpha_{iE}\beta_i)$, then $\mathbf{U}_i(\alpha_i) > \mathbf{U}_i(\beta_i)$.

To see that this condition satisfies B1, note that if $\mathbf{U}_i(\alpha_i) = \mathbf{U}_i(\beta_i) \geq \mathbf{U}_i(\alpha_{iE}\beta_i)$ for all $E \in \Sigma_i$, then it cannot be the case that $\mathbf{U}_i(\alpha_i) > \mathbf{U}_i(\beta_{iE}\alpha_i)$ for $E \in \Sigma_i$. ■

4. Recursive payoffs with multiple priors

In this section, we study a class of games in interim utility form

$$((\Omega_i, \Sigma_i), A_i, \mathbf{V}_i, \mathbf{U}_i)_{i=1}^N ,$$

in which players have multiple priors and interim payoffs satisfying B1 and B2, but not necessarily the conditions in the examples of the previous section. The class includes Bayesian games as a special case.

Player i has a bounded measurable *payoff function*

$$u_i \colon A \times \Omega \to \mathbb{R} ,$$

where the set of *action profiles* A is endowed with the product Borel algebra and $A \times \Omega$ also has the product algebra.

Suppose further that for each i, we are given a set D_i of probability measures on Σ. For each $\pi_i \in D_i$ we write $\hat{\pi}_i$ for its marginal distribution on Σ_i and for the distribution conditional on own types we write $\pi_i(\,\cdot\,|\,\cdot\,) \colon \Sigma_{-i} \times \Omega_i \to [0,1]$, whereby $\pi_i(\,\cdot\,|\omega_i)$ is a probability distribution on Σ_{-i} interpreted as the conditional probability distribution on $\omega_i \in \Omega_i$ realizing.[3]

Assume that the *ex ante* utilities are of the multiple prior form of Gilboa and Schmeidler (1989) and that D_i satisfies the *rectangularity* property of Epstein and Schneider (2003):

$$\begin{aligned} U_i(f) &= \inf_{\pi_i \in D_i} \int_{\Omega} u_i(f(\omega), \omega) \,\mathrm{d}\,\pi_i(\omega) \\ &= \inf_{\pi_i \in D_i} \int_{\Omega_i} \inf_{\nu_i \in D_i} \int_{\Omega_{-i}} u_i(f(\omega), \omega) \,\mathrm{d}\,\nu_i(\omega_{-i}|\omega_i) \,\mathrm{d}\,\hat{\pi}_i(\omega_i) . \end{aligned}$$

With this separation we let

$$\mathbf{V}_i(a_i, f_{-i}|\omega_i) = \inf_{\nu_i \in D_i} \int_{\Omega_{-i}} u_i(a_i, f_{-i}(\omega_{-i}), \omega) \,\mathrm{d}\,\nu_i(\omega_{-i}|\omega_i) ,$$

for each $a_i \in A_i$, $f_{-i} \in F_{-i}$ and $\omega_i \in \Omega_i$. For any interim payoff α_i we let

$$\mathbf{U}_i(\alpha_i) = \inf_{\pi_i \in D_i} \int_{\Omega_i} \alpha_i(\omega_i) \,\mathrm{d}\,\hat{\pi}_i(\omega_i) .$$

We have obtained an associated game in interim utility form and thus a Savage game.

Our first assumption requires that payoffs be continuous in action profiles.

C1. For each $\omega \in \Omega$, the function $a \mapsto u_i(a, \omega)$ is continuous.

We shall also make the following assumption.

C2. There is a probability measure $\mu \colon \Sigma \to [0,1]$ such that

[3] The existence of such a conditional distribution is always guaranteed when the underlying probability space is a Radon space. We note that when each D_i is a singleton and we are in a Bayesian game setting, the existence of equilibrium result in this section does not require the decomposability of priors into marginals and conditionals.

(1) Each $\pi \in \cup_i D_i$ is absolutely continuous with respect to μ.
(2) The marginal distributions of μ over each Σ_i are all atomless.
(3) For each i the set D_i is weak* compact in the dual of $L_\infty(\Omega, \Sigma, \mu)$, the space of real valued μ-essential bounded (equivalence classes) functions on Ω.

In the next assumption we require that the marginal distributions in D_i over Σ_i are mutually absolutely continuous. Epstein and Marinacci (2007) characterize this condition for the maxmin expected utility form in terms of a condition of Kreps (1979).

C3. The marginal densities $\{\hat{\pi}_i \colon \pi_i \in D_i\}$ are mutually absolutely continuous.

Proposition 4.1. *If the game with multiple priors satisfies C1, C2, and C3, then the associated game in interim utility form satisfies B1 and B2. In particular, the associated Savage game satisfies assumptions A1 to A6.*

Proposition 4.1 tells us that when C1, C2, and C3 hold, we need only check that A7 holds to apply Theorem 2.8.

5. Location games with recursive payoffs

We now use this convenient recursive structure to investigate two examples of location games on the sphere. The first is a Bayesian game with payoffs and individual priors that depend on the full profile of types. In the second game, players have multiple priors.

Bayesian location game. Consider the N-player Bayesian game

$$((\Omega_i, \Sigma_i), A_i, u_i, \nu_i)_{i=1}^N ,$$

where (Ω_i, Σ_i), the measurable space of Player i's *types*, is the unit interval $[0, 1]$. The action space A_i of each player is the unit sphere $\mathbb{S}^n$ in $\mathbb{R}^{n+1}$. Player i's prior ν_i is a probability density function $\nu_i \colon \Omega \to \mathbb{R}_+$, which has full support and is Lipschitz continuous.

Let $\mathbb{B}^{n+1}$ denote the unit ball of $\mathbb{R}^{n+1}$. The payoff function $u_i \colon A \times \Omega \to \mathbb{R}$ of Player i is given by

$$u_i(a, \omega) = \gamma_i \left\| P_i(a_i, \omega_i) - R_i(a_{-i}, \omega) \right\|^2 + (1 - \gamma_i) \left\| P_i(a_i, \omega_i) - Q_i(\omega_i) \right\|^2$$

with Lipschitz continuous functions $P_i \colon A_i \times \Omega_i \to \mathbb{S}^n$, $Q_i \colon \Omega_i \to \mathbb{S}^n$, $R_i \colon A_{-i} \times \Omega \to \mathbb{B}^{n+1}$, and $0 \leq \gamma < \frac{1}{2}$. We interpret $R_i(a_{-i}, \omega)$ as Player i's idiosyncratic way of calculating the (generalized) average of the other players' locations, $Q_i(\omega_i)$ as her most preferred location given her type ω_i, and $P_i(a_i, \omega_i)$ as a (possible, but not required) distortion induced by her type ω_i on the degree of her desire to be close to the other players' expected location and her own preferred location. In particular, we allow that Player i may be 'social' for some types, for example, $P_i(a_i, \omega_i) = a_i$, but may be 'anti-social' for other types, for example, $P_i(a_i, \omega_i) = -a_i$. We assume that the inverse correspondence $P_i^{-1} \colon A_i \times \Omega_i \to A_i$ defined by

$$P_i^{-1}(a_i, \omega_i) = \{a_i' \in A_i \colon a_i = P_i(a_i', \omega_i)\} ,$$

is non-empty valued and Lipschitz continuous with constant K. That is, for all $x, y \in A_i \times \Omega_i$ we have

$$\delta\left(P_i^{-1}(x), P_i^{-1}(y)\right) \leq K \|x - y\| ,$$

where δ is the Hausdorff distance between sets in $\mathbb{R}^{n+1}$.

We shall show that this game has a Bayesian Nash equilibrium (in pure strategies). Clearly, B1 and B2 hold. By Proposition 3.1 we need only show that A7 is satisfied.

Fix Player i, a strategy profile f_{-i} of other players, and a type $\omega_i \in \Omega_i$. For each action a_i, let $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ be the *interim expected utility*

$$\mathbf{V}_i(a_i, f_{-i}|\omega_i) = \int_{\Omega_{-i}} u_i(a_i, f_{-i}(\omega), \omega) \, \nu_i(\omega_{-i}|\omega_i) \, \mathrm{d}\, \lambda(\omega_{-i}) ,$$

where λ is the Lebesgue probability measure on $[0,1]^{N-1}$ and

$$\nu_i(\omega_{-i}|\omega_i) = \frac{\nu_i(\omega)}{\int_{\Omega_{-i}} \nu_i(\omega)\,\mathrm{d}\,\lambda(\omega_{-i})}$$

is the conditional probability density of ν_i on Ω_{-i}.

Let

$$M_i(a_{-i},\omega) = \gamma_i R_i(a_{-i},\omega) + (1-\gamma_i) Q_i(\omega_i)\,,$$

and

$$m_i(f_{-i}|\omega_i) = \int_{\Omega_{-i}} M_i(f_{-i}(\omega_{-i}),\omega)\,\nu_i(\omega_{-i}|\omega_i)\,\mathrm{d}\,\lambda(\omega_{-i})\,.$$

We see that

$$\begin{aligned} u_i(a,\omega) &= \|P_i(a_i,\omega_i)\|^2 - 2\langle P_i(a_i,\omega_i), M_i(a_{-i},\omega)\rangle \\ &\quad + \gamma_i \|R_i(a_{-i},\omega)\|^2 + (1-\gamma_i)\|Q_i(\omega_i)\|^2\,, \end{aligned}$$

where for any $x, y \in \mathbb{R}^{n+1}$, $\langle x, y\rangle \in \mathbb{R}$ is the inner-product. Thus

$$\begin{aligned} \mathbf{V}_i(a_i, f_{-i}|\omega_i) &= \|P_i(a_i,\omega_i) - m_i(f_{-i}|\omega_i)\|^2 + \|m_i(f_{-i}|\omega_i)\|^2 \\ &+ \int_{\Omega_{-i}} \gamma_i \|R_i(f_{-i}(\omega_{-i}),\omega)\|^2\,\nu_i(\omega_{-i}|\omega_i)\,\mathrm{d}\,\lambda(\omega_{-i}) + (1-\gamma_i)\|Q_i(\omega_i)\|^2. \end{aligned}$$

Noting that $\|m_i(f_{-i}|\omega_i)\| \geq 1 - 2\gamma_i > 0$, define the point

$$q_i(f_{-i}|\omega_i) = \frac{-m_i(f_{-i}|\omega_i)}{\|m_i(f_{-i}|\omega_i)\|}\,,$$

which is the point on the sphere that is farthest away from $m_i(f_{-i}|\omega_i)$.

Any $a_i \in A_i$ satisfying $q_i(f_{-i}|\omega_i) = P_i(a_i,\omega_i)$, equivalently, $a_i \in P_i^{-1}(q_i(f_{-i}|\omega_i),\omega_i)$, maximizes $\mathbf{V}_i(\cdot, f_{-i}|\omega_{-i})$. In particular, any strategy f_i^* satisfying

$$P_i(f_i^*(\omega_i),\omega_i) = q_i(f_{-i}|\omega_i) \quad \text{equivalently} \quad f_i^*(\omega_i) \in P_i^{-1}(q_i(f_{-i}|\omega_i),\omega_i)$$

for all ω_i is a best response for Player i to f_{-i}.

Now $\omega \mapsto \nu_i(\omega_{-i}|\omega_i)$ is a Lipschitz continuous function because the prior ν_i is a Lipschitz continuous function that is bounded away from zero. Therefore, $\omega_i \mapsto m_i(f_{-i}|\omega_i)$ is also a Lipschitz continuous function with Lipschitz constant K' that is independent of the choice of f_{-i} because of the Lipschitz continuity of Q_i and R_i. This in turn implies that $\omega_i \mapsto q_i(f_{-i}|\omega_i)$ is a Lipschitz continuous function with Lipschitz constant K'', independent of the choice of f_{-i}, because $\gamma_i < \frac{1}{2}$. Finally, we conclude that the closed non-empty valued correspondence $\omega_i \mapsto P_i^{-1}(q_i(f_{-i}|\omega_i),\omega_i)$ is Lipschitz continuous with some constant K^* that is the same for all f_{-i}. By the theorem of Kupka (2005), this correspondence with one dimensional domain has a K^*-Lipschitz continuous selection $\hat{f}_i$, which is a best response to f_{-i}. Letting X_i be the family of K^*-Lipschitz continuous strategies for Player i. By the Arzelà-Ascoli compactness theorem assumption A7 is satisfied.

Location game with multiple-priors. Consider another N-player location game in which once again Player i's type space Ω_i is $[0,1]$ and her action space A_i is the unit sphere $\mathbb{S}^n$. The player's payoff function $u_i\colon A \times \Omega \to \mathbb{R}$ is

$$u_i(a,\omega) = \begin{cases} \|a_i - M_i(a_{-i},\omega)\|^2, & \text{if } \min_{j\neq i}\omega_j \leq \frac{1}{2}\,, \\ 1, & \text{otherwise}\,; \end{cases}$$

where $M_i\colon A_{-i} \times \Omega \to \mathbb{S}^n$ is a Lipschitz continuous function. If the type of at least one of the players other than i is less than or equal to one-half, that is *Low*, then Player i wishes to locate on the circle as far away as possible from $M_i(a_{-i},\omega)$, and may get a payoff greater than one. However, if the type of every player aside from i is greater than a half, that is *High*, then Player i has a guaranteed payoff 1.

We assume the preferences of Player i over strategy profiles take the maxmin expected utility or "multiple priors" form of Gilboa and Schmeidler (1989). For each i, let λ_i be the Lebesgue distribution on Ω_i, and let λ be the product distribution on Ω. Let $\hat{D}_i$ be a weakly compact set of probability density functions on Ω_i in which each $\hat{\nu}_i$ in $\hat{D}_i$ is mutually absolutely continuous with λ_i. Let $\mu_{-i}\colon \Omega_{-i} \times \Omega_i \to \mathbb{R}_+$ and $\nu_{-i}\colon \Omega_{-i} \times \Omega_i \to \mathbb{R}_+$ be functions for which $\omega_i \mapsto \mu_{-i}(\,\cdot\,|\omega_i)$ and $\omega_i \mapsto \nu_{-i}(\,\cdot\,|\omega_i)$ are mappings to conditional probability densities on Ω_{-i}. We assume that for each fixed ω_{-i} the function μ_{-i} is Lipschitz continuous in ω_i. We also assume that if ω_{-i} is in the support of $\mu_{-i}(\cdot|\omega_i)$, then at least one player is of *Low* type. We also assume that for each ω_i, the support of $\nu_{-i}(\,\cdot\,|\omega_i)$ is a subset of $(\frac{1}{2}, 1]^{N-1} \subseteq \Omega_{-i}$.

Now take D_i to be the following set of probability densities defined on Ω:

$$D_i = \{\omega \mapsto \pi_i(\omega_i)(\alpha\mu_{-i}(\omega_{-i}|\omega_i) + (1-\alpha)\nu_{-i}(\omega_{-i}|\omega_i))\colon \pi_i \in \bar{D}_i,\ 0 \le \alpha \le 1\}\,.$$

The *ex ante* utility of Player i for the strategy profile $f \in F$ is

$$U_i(f) = \min_{\pi \in D_i} \int_\Omega u_i(f(\omega), \omega)\pi(\omega)\,\mathrm{d}\,\lambda(\omega)\,.$$

We show that this game has an equilibrium. Since D_i satisfies the *rectangularity* property this is a game in interim form satisfying C1, C2, C3. By Proposition 4.1 we need only show that A7 is satisfied.

Since D_i satisfies the *rectangularity* property it follows if ω_i is realized for Player i then the player wants to maximize the interim utility which in this case is given by

$$\mathbf{V}_i(a_i, f_{-i}|\omega_i) = \min_{\pi_{-i} \in \{\mu_i(\cdot|\omega_i), \nu_i(\cdot|\omega_i)\}} \int_{\Omega_{-i}} u_i(a_i, f_{-i}(\omega_{-i}), \omega)\pi_{-i}(\omega_{-i})\,\mathrm{d}\,\lambda_{-i}(\omega_{-i})\,.$$

For each ω_i and f_{-i} let $m_i(f_{-i}|\omega_i)$ be the point

$$m_i(f_{-i}|\omega_i) = \int_{\Omega_{-i}} M_i(f_{-i}(\omega_{-i}), \omega)\mu_i(\omega_{-i}|\omega_i)\,\mathrm{d}\,\lambda(\omega_{-i})\,.$$

There is a K that is independent of f_{-i} such that $\omega_i \mapsto m_i(f_{-i}|\omega_i)$ is K-Lipschitz continuous.

Fixing ω_i and f_{-i} we notice that for any a_i we have

$$\begin{aligned}
&\int_{\Omega_{-i}} u_i(a_i, f_{-i}(\omega_{-i})\mu_i(\omega_{-i}|\omega_i)\,\mathrm{d}\,\lambda(\omega_{-i}) \\
&\qquad\qquad = \int_{\Omega_{-i}} \|a_i - M_i(f_{-i}(\omega_{-i}), \omega)\|^2 \mu_i(\omega_{-i}|\omega_i)\,\mathrm{d}\,\lambda(\omega_{-i}) \\
&\qquad\qquad \ge 1 \quad = \int_{\Omega_{-i}} u_i(a_i, f_{-i}(\omega_{-i})\nu_i(\omega_{-i}|\omega_i))\,\mathrm{d}\,\lambda(\omega_{-i})
\end{aligned}$$

if and only if

$$\|a_i - m_i(f_{-i}|\omega_i)\|^2 \ge \|m_i(f_{-i}|\omega_i)\|^2.$$

But there is always a point in $\mathbb{S}^n$ satisfying $\|a_i - m_i(f_{-i}|\omega_i)\| \ge \|m_i(f_{-i}|\omega_i)\|$. Therefore, the value $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ of such a point a_i is one. But the maximum of $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ is also one. From this we conclude that a_i maximizes $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ if and only if $\|a_i - m_i(f_{-i}|\omega_i)\| \ge \|m_i(f_{-i}|\omega_i)\|$. In particular, the maximizers of $\mathbf{V}_i$ have the following form:

$$B_i(f_{-i}|\omega_i) = \arg\max_{a_i \in A_i} \mathbf{V}_i(a_i, f_{-i}|\omega_i) = \{a_i \in A_i\colon \langle a_i, m_i(f_{-i}|\omega_i)\rangle \le \tfrac{1}{2}\}\,,$$

for each ω_i and f_{-i}.

This is an upper hemicontinuous correspondence from $[0,1]$ to $\mathbb{S}^n$. That is, there is a K^* independently of f_{-i} such that

$$\delta(B_i(f_{-i}|\omega_i), B_i(f_{-i}|\omega_i')) \le K^*|\omega_i - \omega_i'|\,,$$

for all ω_i, ω_i', where δ is the Hausdorff distance between sets. By the theorem of Kupka (2005), this correspondence has a K^*-Lipschitz continuous selection $\hat{f}_i$. Once again applying the Arzelà-Ascoli theorem yields the desired result.

6. Concluding remarks

We conclude with a discussion of some issues related to the results we have derived in the framework of Savage games.

Universal state space. One question that we do not attempt to answer in this paper is whether it is possible to construct a state space that is a comprehensive representation of the uncertainty faced by players, in the sense of Mertens and Zamir (1985) and Brandenburger and Dekel (1993). We note that Epstein and Wang (1996) do provide such foundations for a setting with purely subjective uncertainty and where the preferences of players need not conform to subjective expected utility theory and so may exhibit non-neutral attitudes toward ambiguity. However, Epstein and Wang's setting does not allow for interdependent preferences. Bergemann, Morris, and Takahashi (2011) construct a universal type space for players with interdependent preferences, but as their framework explicitly involves objective randomization, it is not clear to us how their analysis could be conducted in a Savage setting of purely subjective uncertainty. Finally, Di Tillio (2008) allows for more general preferences, albeit in a setting in which there is only a finite number of outcomes.

Rationalizability. It is also not clear to us what is the appropriate notion of rationalizability in the framework of Savage games. There is an extensive literature that provides foundations for equilibrium in terms of rationalizable behavior, see for example Brandenburger and Dekel (1987) in the context of subjective uncertainty, Tan and Werlang (1988) and Börgers (1993). However, in many of these papers, rationality is expressed in terms of "state-independent expected utility."[4]

To allow for state-dependent ordinal preferences, an alternative notion of rationalizability is needed. As noted by Morris and Takahashi (2012), rationalizability defined in terms of ordinal preferences is invariant to the choice of state space, unlike rationalizability defined in terms of expected utility. However, Morris and Takahashi's notion of rationalizability requires explicit randomization of the kind implied by Anscombe–Aumann acts, which is not available in our setting. Epstein (1997) investigates rationalizability in a setting where strategies are analogs of Savage style acts, nevertheless he rules out state-dependent preferences and restricts the analysis to finite normal form games.

Purification of mixed strategies. We have entirely avoided any assumption on the independence or near independence of player information or types or payoffs. Indeed, in our Bayesian game example, types are statistically dependent via arbitrary Lipschitz-continuous probability density functions. This is in stark contrast with the purification results that follow the classical work of Dvoretzky, Wald, and Wolfowitz (1950), Radner and Rosenthal (1982) and Milgrom and Weber (1985), and related literature.[5] One interpretation of this difference is that while decomposability arguments are also at the heart of purification techniques, those require purification of objectively randomized equilibria. The present paper highlights how our use of the decomposition property can be interpreted as purification of a purely subjective kind.

[4]An exception is Tan and Werlang (1988), who start with a Bayesian game in which players have a state-dependent subjective expected utility function.

[5]Other results include Balder (1988), Khan and Sun (1995), Balder (2002), Khan, Rath, and Sun (2006), Loeb and Sun (2006), Balder (2008), Podczeck (2009), Loeb and Sun (2009), Khan and Rath (2009), Yannelis (2009) and Wang and Zhang (2012).

An important open question is whether it is possible to obtain our results even for standard Bayesian games with interdependent priors using the purification techniques of the extant literature. That literature has focused on the existence of pure-strategy equilibrium in Bayesian games in which information is diffuse. The usual approach is to identify conditions on the information structure of the game that allows us to find a profile of pure strategies that is payoff equivalent to any given equilibrium (randomized) strategy profile. To the best of our knowledge, the techniques that have been developed so far rule out interdependent payoffs and require independent distributions of types.

Games with a separable structure. In general, a Savage game cannot be represented as a Bayesian game. This remains the case, even if there exists an associated game in interim utility form that not only satisfies B1 and B2, but the ex ante utilities of the players are additively separable across states. The difficulty stems from the state-dependence of the players' preferences, which prevents a meaningful separation of beliefs from payoffs (see Karni (1985), Wakker and Zank (1999), and Debreu (1960) for the single decision maker case). Thus it remains an open question as to when can we meaningfully disentangle preferences from beliefs in a Savage game. We outline one approach as our final remark.

Aumann (1974) proposed a class of games with a separable structure to study equilibrium under objective and subjective uncertainty.[6] In this setting, we are able to disentangle subjective beliefs from preferences and thus represent these games as Bayesian games with individual priors. We describe a generalization of the class of games studied by Aumann.[7]

A game with a *separable structure* is an N-player game given by the following tuple:

$$((\Omega_i, \Sigma_i), O_i, \succsim_i, A_i, \zeta_i)_{i=1}^N .$$

Player i is associated with a set of states, Ω_i, and a σ-algebra Σ_i of subsets of Ω_i. She also has an *outcome space* O_i, which we take to be a metric space. The space $\Omega = \times_{i=1}^N \Omega_i$ has the product algebra $\Sigma = \otimes_{i=1}^N \Sigma_i$. An *act* for Player i is a Σ-measurable function $y\colon \Omega \to O_i$. Let Y_i denote the set of Player i's acts. Player i has a preference ordering $\succsim_i$ on the family of acts Y_i. Player i has an action set A_i, which is a compact metric space and a measurable *outcome function* $\zeta_i\colon A \times \Omega \to O_i$, which associates action profiles and state profiles with outcomes. An important difference between this framework and that of Aumann (1974) is that Aumann's outcome function is state independent, that is, it is simply a function from A to O_i.

A *strategy* for Player i is a Σ_i-measurable function $f_i\colon \Omega_i \to A_i$. Let F_i denote the set of Player i's strategies. Each strategy profile $f \in F$ is a Σ-measurable function from Ω to A so there is an induced preference relation $\succsim_i^*$ on F given by

$$f \succsim_i^* g \qquad \text{iff} \qquad \zeta_i \circ f \succsim_i \zeta_i \circ g .$$

The Savage game induced from this game with a separable structure is thus:

$$(\Omega, \bar{A}, (\bar{F}_i, \succsim_i^*)_{i=1}^N) ,$$

where $\bar{A}$ is the disjoint union of the sets A_i, $\bar{F}_i = \{\bar{f}_i\colon f_i \in F_i\}$ for each i, and $\bar{f} \succsim_i^* \bar{g}$ if and only if $f \succsim_i^* g$.

Returning to the game with a separable structure, assume now that for each i there is a probability measure $\pi_i\colon \Sigma \to [0,1]$ and a function $v_i\colon O_i \to \mathbb{R}$ such that

$$V_i(y) = \int_\Omega v_i \circ y(\omega)\, \mathrm{d}\, \pi_i(\omega) ,$$

[6]Aumann and Dreze (2005) develop a related idea in which subjective risk in a game uses available strategies. See also Section 8 of Hammond (2004).

[7]Incidentally, a by-product of such an extension is that, in addition to the standard Bayesian equilibrium notion, Savage games can also be seen to constitute a suitable framework to investigate the existence of (subjectively) correlated equilibrium.

represents $\succsim_i$ over acts for every i. This is the case for example when $\succsim_i$ satisfies all of Savage's postulates and the "monotone continuity" assumption of Arrow (1971) that guarantees that π_i is countably additive. Setting $u_i(a,\omega) := v_i \circ \zeta_i(a,\omega)$ and letting

$$U_i(f) = \int_\Omega u_i(f(\omega),\omega)\,\mathrm{d}\,\pi_i(\omega)$$

for each $f \in F$ we see that U_i is a utility representation of $\succsim_i^*$. We have thus obtained the N-player Bayesian game

$$((\Omega_i,\Sigma_i),A_i,u_i,\pi_i)_{i=1}^N\,.$$

Appendix A. Proofs

Proof of Proposition 2.1. Suppose that $f \sim_i (g_i, f_{-i}) \succsim_i (g_{iE}f_i, f_{-i})$ for every $E \in \mathcal{F}_i$. For any $E \in \mathcal{F}_i$ we have $f \succsim_i (g_{i\Omega\setminus E}f_i, f_{-i})$, thus, by **P2**, $f_{i\Omega\setminus E}g_i, f_{-i} \succsim_i (g_i, f_{-i}) \sim_i f$. □

Proof of Proposition 2.2. It is immediate that $\mathcal{F}_i$ contains $\emptyset$ and Ω. The other two conditions are obtained by noting that

$$g_{iE\setminus E'}f_i = f_{iE'}(g_{iE}f_i)\,,$$

and

$$g_{iE\cup E'}f_i = g_{iE}(g_{iE'}f_i)\,.$$

With this, A3 guarantees that the countable union of events is an event. □

Proof of Corollary 2.3. Let $\sigma(F_i)$ be the smallest σ-algebra of subsets of Ω for which each strategy $f_i \in F_i$ is measurable. Clearly, $\sigma(F_i) \subseteq \Sigma_i \subseteq \mathcal{F}_i$. Pick $E \in \mathcal{F}_i$. Because $|A| \geq 2$, there are $f_i, g_i \in F_i$ such that $E = \{\omega\colon g_{iE}f_i(\omega) \neq f_i(\omega)\}$, which is in $\sigma(F_i)$. Thus, $\Sigma_i = \mathcal{F}_i$ and A3 holds. □

Proof of Proposition 2.4. Clearly, the empty set is in $\mathcal{N}_i$. Let $E \in \mathcal{N}_i$ and $E' \in \mathcal{F}_i$ such that $E' \subseteq E$. If $E' \notin \mathcal{N}_i$, then there are $f \in F$, $g_i \in F_i$ and $j \in N$ satisfying

$$(g_{iE'}f_i, f_{-i}) \not\sim_j f\,.$$

But then

$$((g_{iE'}f_i)_E f_i, f_{-i}) = (g_{iE'}f_i, f_{-i}) \not\sim_j f\,,$$

which is a contradiction, because E is null for Player i.

Furthermore, if $E, E' \in \mathcal{N}_i$, then for any $g_i \in F_i$, by transitivity of $\sim_j$, we have

$$(g_{iE\cup E'}f_i, f_{-i}) = (g_{iE'}(g_{iE}f_i), f_{-i}) \sim_j (g_{iE}f_i, f_{-i}) \sim_j f\,,$$

which tells us that $E \cup E' \in \mathcal{N}_i$. Finally, by A3 and A4, for any increasing sequence of null events E^n, the union E is an event, and it must be null for Player i. □

Proof of Proposition 2.5. We can assume without loss of generality that each $\mathcal{S}_i^m$ also has the the property that if $E', E \in \mathcal{F}$ and $E' \subset E \in \mathcal{S}_i^m$, then $E' \in \mathcal{S}_i^m$.

Fix $\mathcal{S}_i^m$. Denote by $E \ominus E'$ the symmetric difference of any two sets $E, E' \subseteq \Omega$. Let

$$\mathcal{R}_i^m = \{E \in \mathcal{R}_i \colon E \ominus E' \notin \mathcal{N}_i \text{ for all } E' \in \mathcal{S}_i^m\}\,.$$

Proposition A.1. *The following hold true:*

1. *If $E \in \mathcal{R}_i^m$, then $E \notin \mathcal{S}_i^m$.*
2. $\cup_m \mathcal{R}_i^m = \mathcal{R}_i$.
3. *If E^n is an increasing sequence of events whose union E is in $\mathcal{R}_i^m$, then eventually E^n is in $\mathcal{R}_i^m$.*

Proof. (1) is obvious because the empty set is in $\mathcal{N}_i$.

Turning to (2). Suppose that $E \in \mathcal{R}_i$. Suppose by way of contradiction that $E \notin \mathcal{R}_i^m$ for all m. E is associated with $E^m \in \mathcal{S}_i^m$ such that $D^m = E \ominus E^m \in \mathcal{N}_i$. Let $D = \cup_m D^m$, which is in $\mathcal{N}_i$ and we see that $E \setminus D \subseteq E^m$ for all m. Thus, $E \setminus D \in \mathcal{S}_i^m$ for all m. This implies that $E \setminus D$ is in $\mathcal{N}_i$. Thus, $E \in \mathcal{N}_i$, which is a contradiction.

For (3) because $\mathcal{N}_i$ is a σ-ideal, eventually E^n is in $\mathcal{R}_i$. Now if $E^n \notin \mathcal{R}_i^m$ then there exists a null event D such that $E^n \cap D$ is in $\mathcal{S}_i^m$ for all m. By the closedness of $\mathcal{S}_i^m$ $E \cap D \in \mathcal{S}_i^m$ which is impossible. □

If all $\mathcal{R}_i^m$ are empty, then all events are null and the proposition is true trivially. So we can assume that $\mathcal{R}_i^m$ in not empty for all m.

Proposition A.2. *For each m there is $c > 0$ such that*

$$\inf_{E \in \mathcal{N}_i} \max_{\omega \in \Omega} \left(\chi_{\Omega \setminus E} \sum_{k=1}^{n} \alpha^k \chi_{E^k}(\omega) \right) \geq c$$

for any $\alpha^1, \ldots, \alpha^K \geq 0$ and $\sum_{k=1}^{K} \alpha^k = 1$.

Proof. There exists a constant $c > 0$ such that for any finite sequence $E^1, E^2, \ldots, E^n$ in $\mathcal{R}_i^m$ we have

$$\max_{\omega \in \Omega} \frac{1}{n} |\{1 \leq k \leq n \colon \omega \in E^k\}| > c\,.$$

This implies that

$$\max_{\omega \in \Omega} \frac{1}{n} \sum_{k=1}^{n} \chi_{E^k}(\omega) > c$$

where χ_E is the characteristic function of E. This in turn implies that

$$\max_{\omega \in \Omega} \sum_{k=1}^{n} \frac{r^k}{\sum_{k=1}^{n} r^k} \chi_{E^k}(\omega) = \max_{\omega \in \Omega} \frac{1}{\sum_{k=1}^{n} r^k} \sum_{k=1}^{n} r^k \chi_{E^k}(\omega) > c\,,$$

for every $r^k \in \mathbb{N}$, $k = 1, \ldots, n$, satisfying $r^k > 0$ for some k. We conclude that

$$\max_{\omega \in \Omega} \sum_{k=1}^{n} \alpha^k \chi_{E^k}(\omega) \geq c\,,$$

for any convex combination $\alpha^1, \alpha^2, \ldots, \alpha^n \geq 0$, $\sum_{k=1}^{n} \alpha^k = 1$. Therefore,

$$\inf_{E \in \mathcal{N}_i} \max_{\omega \in \Omega} \left(\chi_{\Omega \setminus E} \sum_{k=1}^{n} \alpha^k \chi_{E^k}(\omega) \right) = \inf_{E \in \mathcal{N}_i} \max_{\omega \in \Omega} \sum_{k=1}^{n} \alpha^k \chi_{E^k \setminus E}(\omega) \geq c\,, \tag{1}$$

for any convex combination. □

Let $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ be the ordered vector space of all $\mathcal{N}_i$-equivalence classes of $\mathcal{F}_i$-measurable bounded functions from Ω to $\mathbb{R}$. That is, $f_i \colon \Omega \to \mathbb{R}$ is in $L_\infty(\mathcal{F}_i, \mathcal{N}_i)$ if it is $\mathcal{F}_i$-measurable and bounded, and $g_i \colon \Omega \to \mathbb{R}$ is in the equivalence class $[f_i]$ if $\{\omega \colon f_i(\omega) \neq g_i(\omega)\}$ is in $\mathcal{N}_i$.

For each $f_i \in L_\infty(\mathcal{F}_i|\mathcal{N}_i)$, let

$$\|f_i\|_\infty = \inf_{E \in \mathcal{N}_i} \sup_{\omega \in \Omega} |\chi_{\Omega \setminus E} f_i(\omega)|\,,$$

By Proposition 2.4 $\mathcal{N}_i$ is a σ-ideal of $\mathcal{F}_i$. Thus $\|\cdot\|_\infty$ is a norm on $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$, and with this norm the space is a Banach space. Furthermore, $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ has a canonical ordering whereby $f_i \geq g_i$ if $\{\omega \colon g_i(\omega) > f_i(\omega)\}$ is null. With this vector ordering the Banach space $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ is a Banach lattice with a positive cone $L_\infty^+(\mathcal{F}_i|\mathcal{N}_i)$ that contains any constant function $\mathbf{c} = c\chi_\Omega$, $c > 0$, in its interior.

We list the following result for convenience.

Proposition A.3. *There exists $c > 0$ such that the for any f convex hull C^m in $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ of $\{\chi_E \colon E \in \mathcal{R}_i^m\}$ we have $\|f\|_\infty \geq c$. In particular, C^m is disjoint from $\frac{1}{2}\mathbf{c} - L_\infty^+(\mathcal{F}_i|\mathcal{N}_i)$.*

A separating hyperplane argument now tells us that there is a continuous linear functional π_i^m on $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ separating the two sets. Because zero is an interior point of one set we see that it is non-negative on $L_\infty^+(\mathcal{F}_i|\mathcal{N}_i)$ and that for some $d^m > 0$ we have $\pi_i^m(f_i) > d^m$ for all $f_i \in C^m$.

We can therefore consider π_i^m as a finitely additive measure on $\mathcal{F}_i$. It gives a value of zero to each $E \in \mathcal{N}_i$ and greater than d^m for each $E \in \mathcal{R}_i^m$. By the Hewitt-Yosida decomposition there is a countably additive measure π_i^{mc} and a purely finitely additive measure π_i^{mf} such that

$$\pi_i^m = \pi_i^{mc} + \pi_i^{mf}\,.$$

Pick $E \in \mathcal{R}_i^m$. Because π_i^{mf} is purely finitely additive, for $d^m > \alpha > 0$ there is an increasing sequence $E^n \in \mathcal{F}_i$, $\cup_n E^n = E$, and

$$\lim_n \pi_i^{mf}(E^n) \leq \alpha\,.$$

But $E^n \in \mathcal{R}_i^m$ eventually for n large enough. From this we conclude that for such n

$$\pi_i^{mc}(E) \geq \pi_i^{mc}(E^n) = \pi_i^m(E^n) - \pi_i^{mf}(E^n) \geq \gamma^m - \alpha > 0\,.$$

Normalize each π_i^{mc} making it a probability measure and consider the probability measure:

$$\pi_i = \frac{1}{2^m}\sum_{m=1}^{\infty} \pi_i^{mc}(E)\,.$$

We see that $\pi_i(E) > 0$ for each $E \in \mathcal{R} = \cup_m \mathcal{R}_i^m$. That is, π_i is the required measure. □

Proof of Proposition 2.6. Suppose that $E \in \mathcal{R}_i$. For some player $j \in N$, some strategy profile $f \in F$, and some strategy $g_i \in F_i$ we have $(g_{iE} f_i, f_{-i}) \not\sim_j f$. Thus, there exists a sequence $\{E^1, \ldots, E^k\} \subseteq \mathcal{F}_i$, $\cup_k E^k = \Omega$, satisfying $(f_{iE^k}(g_{iE} f_i), f_{-i}) \not\sim_j f$ for all k. But $f_{iE^k}(g_{iE} f_i) = g_{iE \setminus E^k} f_i$. Thus, for all k the event $E \setminus E^k$ is in $\mathcal{R}_i$. Also, because $\mathcal{N}_i$ is an ideal there must be some k^* such that $E^{k^*} \cap E$ is relevant for Player i. The relevant events $E \setminus E^{k^*}$ and $E^{k^*} \cap E$ are disjoint and their union is E. □

Proof of Proposition 2.7. We need only show that (2) implies (1). Let π_i be from condition (2). If $\pi_i(\Omega) = 0$, then we are done, since there are no relevant events. Otherwise, without loss of generality we can assume $\pi_i(\Omega) = 1$.

For each m, define the set

$$\mathcal{S}_i^m = \{E \in \mathcal{F}_i : \pi_i(E) \leq 1/m\}\,.$$

Notice that each $\mathcal{S}_i^m$ is closed and note that $\cap_m \mathcal{S}_i^m \subseteq \mathcal{N}_i$.

Now let $E^1, E^2, \ldots, E^n$ is a finite sequence in $\mathcal{F}_i$ not in $\mathcal{S}_i^m$. We have

$$\begin{aligned}
\max_{\omega \in \Omega} \frac{1}{n} |\{1 \leq k \colon \omega_i \in E^k\}| &= \frac{1}{n} \max_{\omega \in \Omega} \sum_{k=1}^{n} \chi_{E^k} \\
&\geq \frac{1}{n} \int_\Omega \sum_{k=1}^{n} \chi_{E^k}(\omega)\, \mathrm{d}\, \pi_i(\omega) \\
&= \frac{1}{n} \sum_{k=1}^{n} \pi_i(E^k) \\
&\geq \frac{1}{n}(n\frac{1}{m}) = \frac{1}{m}\,.
\end{aligned}$$

This shows that A5 holds.

We show that A6 also holds. Suppose that $(g_i, f_{-i}) \not\sim_j f$. Suppose by way of contradiction that for each m there is E^m satisfying $\pi_i(E^m) \le \frac{1}{m}$ such that $(h_{iE^m} g_i, f_{-i}) \sim_j f$. Noting that χ_{E^m} converges to zero in π_i-measure we can move to a subsequence such that χ_{E^m} converges π_i-almost surely to zero. But zeros of π_i are all null for Player i. Thus, by A4 $(g_i, f_{-i}) \sim_j f$, which is a contradiction. □

A fixed point theorem. We begin with a statement of a fixed point theorem and apply it to prove Theorem 2.8. For a complete proof of the fixed point theorem used in this subsection, please refer to Meneghel and Tourky (2013).

Let (S, Σ, μ) be an atomless probability space and let T be a topological space. Let $L(S, T)$ be the set of all functions, not necessarily measurable, from S to T. Endow $L(S, T)$ with the topology of pointwise convergence.

A set-valued (possibly empty valued) mapping $B\colon F \twoheadrightarrow F$ is a *decomposable mapping* if its domain F and values $B(f)$, for all $f \in F$, are decomposable sets. A decomposable mapping B is *μ-sequentially closed graphed* if the following hold:

(1) If $\mu(E) = 0$ and $g \in B(f)$, then $h_E g \in B(f)$ and $g \in B(h_E f)$ for all $h \in F$.
(2) F is sequentially closed in $L(S, T)$.
(3) B has a sequentially closed graph in $F \times F$.

A *fixed point* of B is a function $f \in F$ satisfying $f \in B(f)$.

Theorem A.4 (Corollary 2.3, Meneghel and Tourky (2013))**.** Let $B\colon F \twoheadrightarrow F$ be a decomposable μ-sequentially closed-graphed mapping. If for a compact and metrizable $X \subseteq F$ we have $X \cap B(f) \neq \emptyset$ for each $f \in F$, then B has a fixed point.

Proof of Theorem 2.8. Assume that A7 holds.

For each i let $S_i = \Omega$ and $T_i = A$. Each $f \in F$ is can be considered a function from S to T whereby

$$f(s_1, \dots, s_N) = (f_1(s_1), \dots, f_N(s_N))\,.$$

For each i consider the atomless measure space $(S_i, \mathcal{F}_i, \pi_i)$ from Proposition 2.7. We will assume that at least one player has a relevant event and that all non-zero π_i are probability measures. Let $\mathcal{F} = \otimes_{i=1}^N \mathcal{F}_i$ be the tensor product. Each E in $\mathcal{F}$ that is not the empty set is of the form

$$(E_1, E_2, \dots, E_N)\,,$$

where $E_i \in \mathcal{F}_i$ for each i. Now if $f, g \in F$, then

$$g_E f(s_1, \dots, s_N) = (g_{1E_1} f_1(s_1), \dots, g_{NE_N} f_N(s_n))\,,$$

which is in F.

Let $\mu\colon \mathcal{F} \to [0, 1]$ be the probability measure given by

$$\mu(E) = \frac{1}{N} \sum_{i=1}^{N} \pi_i(E_i)\,.$$

This, is an atomless measure and if $\mu(E) = 0$, then each E_i is null for Player i. For each $f \in F$, let

$$B(f) = \{g \in F\colon g_i \text{ is a best response to } f_{-i} \text{ for all } i\}\,.$$

Notice that if $\mu(E) = 0$ and $g \in B(f)$ we have $h_E g \in B(f)$ and $g \in B(h_E f)$.

Now our sets X_i are compact and metrizable in the topology of pointwise convergence. Therefore, their product $X \subseteq F$ is compact and metrizable in the same topology. Assume first that there is only one player. Clearly, an equilibrium exists because the player maximizes her preferences in the compact and metrizable set X. Now suppose that there are two or more players. By assumption A7, the sequentially closed, decomposable set $\widetilde{X}_i$ is a subset of F_i for each i. Let $\widetilde{X}$ be the product of X_i, which is sequentially closed and decomposable once again. Let $\widetilde{B}\colon \widetilde{X} \to \widetilde{X}$ be the restriction of B to $\widetilde{X}$. We see that B

is a decomposable mapping that is also μ-sequentially closed graphed. Applying Theorem A.4 gives us the required equilibrium.

For the converse, if f is an equilibrium, then letting $X_i = \{f_i\}$ for each i and noting that because single points in A are closed, we have that $\widetilde{X}_i = X_i$, A7 holds. □

Proof of Proposition 3.1. Suppose that $f \in F$ and $g_i \in F_i$ satisfy

$$U_i(f) = U_i(g) \geq U_i(g_{iE_i} f_i, f_{-i})$$

for all $E_i \in \Sigma_i$. By B1 this means that

$$U_i(g_{iE_i} f_i, f_{-i}) = \mathbf{U}_i(\mathbf{V}_i(g_{iE_i} f_i, f_{-i})) = \mathbf{U}_i(\mathbf{V}_i(g_i, f_{-i})_{E_i} \mathbf{V}_i(f)) = \mathbf{U}_i(\mathbf{V}_i(f)) = U_i(f)$$

$E_i \in \Sigma_i$. Thus, A2 holds.

That A3 is satisfied is a consequence of Corollary 2.3.

Now U_i is continuous for pointwise convergent sequences of strategy profiles by B2. So A4 holds.

If A has less than two points, then all events are null and A5 and A6 hold trivially. If they have two or more points, then by Corollary 2.3 $\bar{\Sigma}_i = \mathcal{F}_i$. Now the restriction of μ_i to $\bar{\Sigma}_i$ is atomless and if $\mu(E) = 0$, then E is null for Player i by (2) of B2. By Proposition 2.7 assumptions A5 and A6 hold. □

Proof of Proposition 4.1. Consider the associated game in interim form. Clearly, B2 holds. For B1 let α_i, β_i be the two interim payoffs. Choose $\mu_i^* \in D_i$ in

$$\arg\min_{\mu_i \in D_i} \int_{\Omega_i} \alpha_i \vee \beta_i \,\mathrm{d}\,\hat{\mu}_i(\omega_i)\,.$$

We see that

$$\mathbf{U}_i(\alpha_i) \geq \mathbf{U}_i(\alpha_i \vee \beta_i) \geq \int_{\Omega_i} \alpha_i(\omega_i) \,\mathrm{d}\,\hat{\mu}_i^*(\omega_i) \geq \mathbf{U}_i(\alpha_i)\,.$$

Thus, it must be the case that α_i and $\alpha_i \vee \beta_i$ agree μ_i^*-almost surely. Similarly, β_i and $\alpha_i \vee \beta_i$ agree μ_i^*-almost surely. This implies that α_i and β_i agree almost surely for all $\mu_i \in D_i$. This implies that $\mathbf{U}_i(\beta_{iE}\alpha_i) = \mathbf{U}_i(\alpha_i)$ for all $E_i \in \Sigma_i$, as required. □

References

Arrow, K. J. (1971). *Essays in the theory of risk-bearing*, Volume 1. Markham Publishing Company Chicago.

Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica 69*(4), 861–889.

Aumann, R. and J. Dreze (2005). Assessing strategic risk. *Available at SSRN 880392*.

Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics 1*(1), 67–96.

Balder, E. J. (1988). Generalized equilibrium results for games with incomplete information. *Mathematics of Operations Research 13*(2), 265–276.

Balder, E. J. (2002). A unifying pair of CournotNash equilibrium existence results. *Journal of Economic Theory 102*(2), 437–470.

Balder, E. J. (2008). Comments on purification in continuum games. *International Journal of Game Theory 37*(1), 73–92.

Bergemann, D., S. Morris, and S. Takahashi (2011). Interdependent preferences and strategic distinguishability. Princeton University Economic Theory Research Paper No. 008-2011.

Börgers, T. (1993). Pure strategy dominance. *Econometrica 61*(2), 423–430.

Brandenburger, A. and E. Dekel (1987). Rationalizability and correlated equilibria. *Econometrica 55*(6), 1391–1402.

Brandenburger, A. and E. Dekel (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory 59*(1), 189–198.

Debreu, G. (1960). Topological methods in cardinal utility theory. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical methods in the social sciences*, pp. 16–26. Stanford CA: Stanford University Press.

Denneberg, D. (1994). *Non-additive measure and integral*, Volume 27. Kluwer Academic Pub.

Di Tillio, A. (2008). Subjective expected utility in games. *Theoretical Economics 3*(3), 287–323.

Dvoretzky, A., A. Wald, and J. Wolfowitz (1950). Elimination of randomization in certain problems of statistics and of the theory of games. *Proceedings of the National Academy of Sciences of the United States 36*(4), 256–260.

Epstein, L. G. (1997). Preference, rationalizability and equilibrium. *Journal of Economic Theory 73*, 1–29.

Epstein, L. G. and M. Marinacci (2007). Mutual absolute continuity of multiple priors. *Journal of Economic Theory 137*, 716–720.

Epstein, L. G. and M. Schneider (2003). Recursive multiple-priors. *Journal of Economic Theory 113*, 1–31.

Epstein, L. G. and T. Wang (1996). "Beliefs about beliefs" without probabilities. *Econometrica 64*(6), 1343–1373.

Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics 16*(1), 65–88.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Grant, S., A. Kajii, and B. Polak (2000). Decomposable choice under uncertainty. *Journal of Economic Theory 92*(2), 169–197.

Hammond, P. J. (2004). Expected utility in non-cooperative game theory. *Handbook of Utility Theory: Volume 2: Extensions 2*, 979.

Karni, E. (1985). *Decision making under uncertainty: the case of state-dependent preferences*. Cambridge, MA: Harvard University Press.

Karni, E., D. Schmeidler, and K. Vind (1983). On state dependent preferences and subjective probabilities. *Econometrica 51*(4), 1021–1031.

Kelley, J. L. (1959). Measures on Boolean algebras. *Pacific Journal of Mathematics 9*(11), 1165–1177.

Khan, M. A. and K. P. Rath (2009). On games with incomplete information and the Dvoretsky–Wald–Wolfowitz theorem with countable partitions. *Journal of Mathematical Economics 45*(12), 830–837.

Khan, M. A., K. P. Rath, and Y. Sun (2006). The Dvoretzky-Wald-Wolfowitz theorem and purification in atomless finite-action games. *International Journal of Game Theory 34*(1), 91–104.

Khan, M. A. and Y. Sun (1995). Pure strategies in games with private information. *Journal of Mathematical Economics 24*(7), 633–653.

Kreps, D. (1979). A representation theorem for preference for flexibility. *Econometrica 47*, 565–577.

Kupka, I. (2005). Continuous selections for Lipschitz multifunctions. *Acta Mathematica Universitatis Comenianae 74*(1), 133–141.

Loeb, P. and Y. Sun (2006). Purification of measure-valued maps. *Illinois Journal of Mathematics 50*(1-4), 747–762.

Loeb, P. and Y. Sun (2009). Purification and saturation. *Proceedings of the American Mathematical Society 137*(8), 2719–2724.

McAdams, D. (2003). Isotone equilibrium in games of incomplete information. *Econometrica 71*(4), 1191–1214.

Meneghel, I. and R. Tourky (2013). A fixed point theorem for closed-graphed decomposable-valued correspondences. Mimeo.

Mertens, J.-F. and S. Zamir (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory 14*(1), 1–29.

Milgrom, P. R. and R. J. Weber (1985). Distributional strategies for games with incomplete information. *Mathematics of Operations Research 10*(4), 619–632.

Morris, S. and S. Takahashi (2012). Games in preference form and preference rationalizability. Princeton University Economic Theory Research Paper No. 43-2012.

Podczeck, K. (2009). On purification of measure-valued maps. *Economic Theory 38*(2), 399–418.

Radner, R. and R. W. Rosenthal (1982). Private information and pure-strategy equilibria. *Mathematics of Operations Research 7*(3), 401–409.

Reny, P. J. (2011). On the existence of monotone pure strategy equilibria in Bayesian games. *Econometrica 79*(2), 499–553.

Rockafellar, R. (1968). Integrals which are convex functionals. *Pacific Journal of Mathematics 24*(3), 525–539.

Rockafellar, R. (1971). Integrals which are convex functionals, II. *Pacific Journal of Mathematics 39*(2), 439–469.

Savage, L. J. (1954). *The foundations of statistics.* John Wiley & Sons.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*(3), 571–587.

Tan, T. C.-C. and S. R. D. C. Werlang (1988). The Bayesian foundations of solution concepts of games. *Journal of Economic Theory 45*, 370–391.

Wakker, P. and H. Zank (1999). State dependent expected utility for Savage's state space. *Mathematics of Operations Research 24*(1), 8–34.

Wang, J. and Y. Zhang (2012). Purification, saturation and the exact law of large numbers. *Economic Theory 50*(3), 527–545.

Yannelis, N. C. (2009). Debreus social equilibrium theorem with asymmetric information and a continuum of agents. *Economic Theory 38*(2), 419–432.

School of Economics, The University of Queensland, St Lucia QLD 4072, Australia.
E-mail address: simongrant@uq.edu.au
E-mail address: i.meneghel@uq.edu.au
E-mail address: rabee.tourky@gmail.com