Theoretical Economics 11 (2016), 641–682 1555-7561/20160641

Savage games

Simon Grant
Research School of Economics, Australian National University and
Department of Economics, The University of Queensland

Idione Meneghel
Research School of Economics, Australian National University

Rabee Tourky
Research School of Economics, Australian National University

We define and discuss Savage games, which are ordinal games of incomplete information 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 class of games we consider, player preferences satisfy versions of Savage's *sure-thing principle* and *small event continuity* postulate. Savage games provide a tractable framework for studying attitudes toward uncertainty in a strategic setting. The work eschews any notion of objective randomization, convexity, monotonicity, or independence of beliefs. We provide a number of examples illustrating the usefulness of the framework, including novel results for a purely ordinal matching game that satisfies all of our assumptions and for games for which the preferences of the players admit representations from a wide class of decision-theoretic models.

Keywords. Subjective uncertainty, strategic interaction, strategically irrelevant events, ambiguity, Bayesian games.

JEL classification. C72, D81.

1. Introduction

Harsanyi (1967–1968) introduced the class of Bayesian games to analyze games of incomplete information. These are strategic interactions in which one or more of the participants is uncertain about some relevant aspect of the game being played. As Harsanyi

Simon Grant: simon.grant@anu.edu.au
Idione Meneghel: idione.meneghel@anu.edu.au
Rabee Tourky: rabee.tourky@anu.edu.au
We are especially grateful for the insightful comments and suggestions from Dirk Bergemann, Larry Epstein, Atsushi Kajii, Edi Karni, Ali M. Khan, Stephen Morris, and Larry Samuelson. We also thank Jean-Marc Bonnisseau, Michael Greinecker, Alejandro Jofre, George Mailath, Clemens Puppe, Luca Rigotti, Roee Teper, and Nicholas Yannelis for helpful comments. The paper has greatly benefited from careful reports of referees. In particular, we would like to thank the referees for help in improving the motivations in the paper.

DOI: 10.3982/TE2068

noted, participants might not know precisely what consequence will result from the actions chosen, or they may be unsure about what actions are actually available to the other participants, or they may be uncertain about the other participants' preferences. His key insight was to assume that each participant deals with her incomplete information by assigning a *subjective* probability distribution over all the variables not known to her and chooses an action out of the set of those available to her that maximizes the mathematical expectation of her payoff in terms of that probability distribution. In such a formulation, however, the preferences of the participants rather than being primitives of the model are instead *constructed* from the (cardinal) subjective probabilities and payoffs. Moreover, such a subjective expected utility representation necessarily rules out behavior that can accommodate nonneutral attitudes toward ambiguity.

In this paper, we suggest an alternative way to model the strategic interaction of a finite number of players in the presence of purely subjective uncertainty. To do so, we adopt the framework of Savages's theory of decision-making under uncertainty, and adapt it to allow for strategic interaction among the players. In the process, we introduce a new class of games, denoted *Savage games*, in which the players' objects of choice (strategies) are identified with Savage acts (state-contingent plans), and we propose a solution concept that captures the idea of players lacking a strict incentive to deviate at equilibrium. The following desiderata guide our design of a choice-based theory of strategic interaction under subjective uncertainty:

1. *Expressed in terms of preferences.* Equilibrium behavior seems more basic than any particular functional form representing utilities. Therefore, we follow the decision-theoretic tradition and write our assumptions in terms of basic preferences.

2. *Consistency with the Savage framework.* In the single-player setting the class of games reduces to Savages's framework for individual decision-making under uncertainty. Hence we can include generalizations of subjective expected utility such as those allowing for nonneutral attitudes toward ambiguity, thereby allowing for the incorporation of such attitudes in multiplayer settings.

3. *Consistency with Bayesian games.* The class of games includes Bayesian games as a proper subclass. Moreover, within that subclass the solution is equivalent to Bayesian equilibrium.

4. *Theoretical consistency.* The framework is sufficiently rich to accommodate assumptions that guarantee the existence of equilibrium. In the sequel we show that these assumptions are closely related to the assumptions made in models of subjective expected utility and its generalizations.

5. *Parsimony.* The framework and the assumptions imposed therein are based on the smallest set of elements needed to specify the strategic interaction in which the participants facing subjective uncertainty are engaged.

With all of this in mind, our purpose in this paper is to consider a wider circle of issues than simply the characterization of a class of preferences that admit a particular form of representation. Similarly, our purpose is not to provide an epistemic foundation for decision-making in the presence of subjective strategic uncertainty.

Savages's framework provides a natural starting point for the study of strategic interactions among players in the presence of subjective uncertainty. It allows us, as well as behooves us, to model equilibrium behavior without the usual technical paraphernalia of convexity or monotonicity of strategies and preferences, and the related praxis that seems to have arisen more from considerations of analytical tractability rather than motivated by, for example, behavioral properties of the underlying preferences.

Going beyond existence of equilibrium, the Savage games framework potentially fills a gap in both the decision-theoretic literature and game theory by providing a purely subjective architecture for understanding attitudes toward uncertainty and risk where the uncertainty a decision-maker faces is generated, in part, by the behavior of other individuals. For example, it provides a way to begin to understand incentives by individuals to exploit attitudes toward ambiguity of other individuals: all of this is articulated in a purely subjective framework unlike Azrieli and Teper (2011), Bade (2011), and Riedel and Sass (2014).

Working with such a framework, however, poses a major challenge as it precludes the use of most of the techniques available in the extant literature on equilibrium theory. First, in a setting that does not involve any notion of objective randomization, the lack of convexity rules out classical calculus approaches as well as the geometric analysis developed over the past century in economic theory. Second, without a natural order structure and corresponding intrinsic notion of monotonicity, it is not possible to use the more recent order-theoretic ideas and associated results, as in Reny (2011). Without delving into the technical details, let us just note at this juncture, that our theorem for the presence of an equilibrium is crafted around the standard consequentialist reasoning embodied in Savage's *sure-thing principle* exploiting the ability to move from one best response to another by means of the decomposable choice property of the preferences.

At a first glance, the way our assumptions are stated prompts the comparison between our model and standard axiomatic theories of individual choice. This immediately raises the issue of revealed preferences. We have (partially) addressed the question of observability by making assumptions only on the preferences of each player when keeping fixed the strategy profile of the other players. As in any game-theoretic setting, however, the presence of strategic interaction aggravates the problem of observability of preferences: one may always claim that only equilibrium choices are actually observable.

But we make no claim that our framework provides a solution to the problem of observability of preferences in a strategic setting. Indeed, our assumptions, although related, do not play the same role as the usual axioms in utility representation results. Rather, they are only sufficient conditions on the ordinal ranking of alternatives faced by players to guarantee the existence of a strategy profile secure against profitable deviations. Our goal is to predict the outcome of a situation where there is strategic interaction in the presence of subjective uncertainty. And since the equilibrium concept is ordinal in nature, by making assumptions on the players' ordinal ranking of alternatives, we are able to establish existence without resorting to any notion of objective randomization or convexity.

Although relevant, the question of observability of preferences in a strategic setting is outside the scope of this paper. It is an important question regarding the decision-theoretic foundations of game theory, and we refer the reader to the following references that have attempted to tackle this difficult problem: Fishburn (1982), Gilboa and Schmeidler (2003), and Aumann and Dreze (2008).

The paper is organized as follows. First, by means of a simple example, we illustrate in Section 2 the environment and the main ideas in the model. In Section 3 we describe and study Savage games, introducing and motivating the assumptions that underpin our main theorem on the existence of equilibrium. In Section 4 we provide an example of a purely ordinal matching game that satisfies all the assumptions required for our existence result, thereby illustrating the notion of ordinal equilibrium. In Section 5 we study games with *recursive* payoffs. 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 6 contains two examples with recursive payoffs, the first a Bayesian game and the second with (recursive) multiple priors. We conclude in Section 7 with a discussion of related work, open questions, and possible extensions of this work. In particular, we highlight a possible approach to a major unresolved question regarding the axiomatic characterization of Bayesian games within the Savage games framework.

All the proofs can be found in the Appendix.

2. The environment and an example

Our aim in this paper is to develop an ordinal class of games for analyzing strategic interaction involving incomplete information with purely subjective uncertainty that eschews any a priori specification of the cardinal aspects of Bayesian games such as utilities and/or probabilities. Instead any intrinsic attitudes toward uncertainty will be identified (to the extent that this is possible) solely from the (ordinal) preferences of the participants that guide the choices they make in the “play” of the game. Thus everything that any of the participants may be uncertain about is encoded in a *state space* that comprises a collection of mutually exclusive and exhaustive *states* (of the world). As in Savage (1954), each state will be taken to be “a description of the world, leaving no relevant aspect undescribed” (Savage 1954, p. 9). Although Savage was providing a foundation for a theory of individual choice under uncertainty, he viewed his approach as one that could accommodate considerable generality in terms of the degree of specificity and comprehensiveness of the description of a state of the world. Indeed one of his (motivating) examples of what such a description could conceivably entail was “... [t]he exact and entire past, present, and future history of the universe, understood in any sense, however wide” (Savage 1954, p. 8). For any particular strategic interaction, however, we concur with Savages’s advocacy for “the use of modest little worlds, tailored to particular contexts...” (Savage 1954, p. 9). Thus the formal specification of the state space need only be as rich as is required to accommodate the universe of things that the participants are not sure about and that are *relevant* to the specific strategic interaction being analyzed.

We define a strategy for a player to be a *state-contingent* choice of action. We refer to the collection of strategies available to that player as her strategy set. A consequence will then be anything that may happen as a result of the actions chosen by the participants and the state of the world that obtains. One significant feature of strategic interaction under uncertainty that is not present in Savages's framework for individual choice under uncertainty is that, in general, different participants will have different information about which state of the world obtains. We shall interpret the information available to a player as corresponding to those subsets of the state space over which she can "condition" her choice of strategy. That is, we shall interpret any subset of the state space as an *(information) event* for this player if she is able to "deviate" from any of her available strategies on that event to any other of her available strategies.

To illustrate these ideas, consider the following strategically interactive elaboration of Savages's omelet example (Savage 1954, pp. 13–15). Leonard has just broken five good eggs into a bowl when the front doorbell to his apartment sounds. Jimmie his roommate comes into the kitchen and volunteers to finish making the omelet and to clean up, allowing Leonard to go and open the door, which is not visible from the kitchen and is far enough away from the kitchen that anything said there cannot be heard by anyone in the kitchen. They both think the caller could be either their landlord or their neighbor, Jane, on whom, everybody knows, Leonard has a crush. Moreover, everyone knows that the feelings he has toward Jane are reciprocated by her. Leonard is unsure, however, about Jimmie's attitude toward this budding romance. In particular, he does not know whether Jimmie is envious or happy for them. If the caller turns out to be Jane, then Leonard must decide whether or not to invite her to come in to share their omelet.[1] A sixth egg, intended for the omelet, lies unbroken beside the bowl containing the five good eggs. Jimmie must decide either to break this egg into the bowl containing the five good eggs or to break it into a saucer affording him the opportunity to check whether it is rotten before adding it to the other five good eggs.

For some reason Leonard and Jimmie (must) make their respective decisions simultaneously. To ensure that it is only Leonard and Jimmie who have a nontrivial decision to make, let us assume that Jane will accept for sure any invitation to come in and share their omelet.

In this description of strategic interaction under uncertainty in which Leonard and Jimmie are the two players, there are three things that either one or both of them are uncertain about:

(i) whether or not the egg is rotten

(ii) whether the caller at the door is their landlord or is Jane

(iii) whether Jimmie is envious of or is happy about Leonard and Jane's romance.

[1] We suppose that their relationship with their landlord although proper and polite is not one that could be said to be "familiar" to any meaningful degree and so neither would ever consider it appropriate to invite their landlord in to share a meal.

As these three things are at least logically distinct, it is natural to model the uncertainty with a state space that has 2^3 elements. For example, we could specify the state space as

$$\Omega = \{0, 1\} \times \{0, 1\} \times \{0, 1\},$$

where each $\omega = rst \in \Omega$ is the state in which

$$r = \begin{cases} 1 & \text{if the sixth egg is rotten} \\ 0 & \text{if the sixth egg is good} \end{cases}$$

$$s = \begin{cases} 1 & \text{if the caller at the door is the landlord} \\ 0 & \text{if the caller at the door is Jane} \end{cases}$$

$$t = \begin{cases} 1 & \text{if Jimmie is envious of Leonard and Jane's romance} \\ 0 & \text{if Jimmie is happy about Leonard and Jane's romance.} \end{cases}$$

From the description above we can also identify four distinct actions that may be taken in the course of "play." So, for example, we could take the action set A to be the four-element set $\{a_1, a_2, a_3, a_4\}$, where action

a_1 is breaking the sixth egg into the bowl containing the other five good eggs

a_2 is breaking the sixth egg into a saucer for inspection

a_3 is inviting the caller at the door in to share the omelet

a_4 is not inviting the caller at the door in to share the omelet.

A consequence can then be associated with each (feasible) action pair and state combination. For example, the action pair and state combination $(a_1, a_4, 111)$ is associated with the consequence that results when Jimmie, who is envious of Leonard and Jane's romance, breaks a rotten egg into the bowl containing the other five good eggs. This saves Jimmie the bother of having to wash an extra saucer, but in this instance at the cost of ruining the omelet. Moreover, when Leonard opens the door, he finds it is their landlord, to whom he does not extend an invitation to come in to join them and share their omelet.

We take the strategies available to our players, Leonard and Jimmie, to be those mappings from the state space Ω to the action set A that are *consonant* with the above description of their strategic interaction under uncertainty in terms of what each player knows and does not know and consequently what actions are available to him. For example, if f_J denotes a strategy available to Jimmie, then its range must be a nonempty subset of $\{a_1, a_2\}$, since Jimmie does not go to answer the door and so cannot invite Jane to share the omelet, should she be the one who was ringing their front doorbell. Moreover, assuming that Jimmie knows his own feelings about Leonard and Jane's romance, then the only aspect of the state $\omega = rst$ on which he can condition his decision whether to select action a_1 or action a_2 is the t. Notice, when he is deciding between breaking the sixth egg into the bowl with the five good eggs or breaking it into a separate saucer, he does not know whether the egg is rotten or good; neither does he know who is the caller at the door.

If we denote the set of strategies available to Jimmie by F_J, it follows from the above description of the strategic interaction that

$$F_J = \{f_J : \Omega \to \{a_1, a_2\} \mid f(00t) = f(01t) = f(10t) = f(11t), t = 0, 1\}.$$

Correspondingly, if f_L denotes an available strategy for Leonard, then it follows from the description of the situation above that, for any state $\omega = rst$ with $s = 1$ (that is, the event in which the caller at the door is the landlord), $f_K(\omega) = a_4$, since we have assumed, given the nature of their relationship with their landlord they would never contemplate inviting him in to share a meal.[2] So, if we denote the set of strategies available to Leonard by F_L, then it follows from our story above that

$$F_L = \{f_L : \Omega \to \{a_3, a_4\} \mid f(111) = a_4 \text{ and } f(0s0) = f(0s1) = f(1s0) = f(1s1), s = 0, 1\}.$$

Notice that we can associate with each strategy profile (f_J, f_L) in $F_J \times F_L$ the following mapping from states to consequences:

$$\omega \mapsto (f_J(\omega), f_L(\omega), \omega).$$

Thus, from the perspective of the analyst or modeler, the ex ante uncertainty facing the players given they collectively choose their actions according to that strategy profile is embodied in this associated *state-contingent consequence*. It is therefore natural to assume that each player's choice of strategy will be guided by her underlying preferences over state-contingent consequences that will be reflected by a binary relation defined over the set of strategy profiles. Hence, to complete the specification of the ordinal normal-form game for modeling Leonard and Jimmie's strategic interaction under uncertainty, requires two binary relations $\succsim_J, \succsim_L \subset F_J \times F_L$ corresponding to the (ex ante) preferences of Jimmie and Leonard, respectively, that will guide their strategy choice.

An equilibrium for this ordinal normal-form game is a strategy profile (f_J^*, f_L^*) for which

$$(f_J^*, f_L^*) \succsim_J (f_J, f_L^*) \quad \text{for all } f_J \in F_J \quad \text{and}$$
$$(f_J^*, f_L^*) \succsim_L (f_J^*, f_L) \quad \text{for all } f_L \in F_L.$$

Notice that for the purpose of finding such an equilibrium (should one exist) there is no need for one player to be able to express a preference between two strategy profiles that involve *different* strategy choices by the other player. Although such a preference may be plausible from the description of the strategic interaction, we note that it involves comparisons across pairs of strategy profiles that cannot inform us about how

[2]Alternatively we could consider expanding Leonard's strategy set to include strategies that involve him extending an invitation to the landlord and through the specification of his preferences ensure that it would never be a best response for him to select any strategy that involved him inviting the landlord in. However, unless we presumed a predetermined response by the landlord this would mean the landlord would become a third strategic player whose behavior we would need to consider. Moreover, our purpose here is to illustrate how the uncertainty one player may have about the actions available to another player can be modeled within the Savage game framework.

that player might choose her strategy in the ordinal game. For example, for any given strategy of Leonard's $f_L \in F_L$, if Leonard had the strict preference $(f_J, f_L) \succ_L (f'_J, f_L)$ (where $\succ_L$ denotes the asymmetric component of $\succsim_L$), then no matter what preferences Jimmie had, this strict preference could never guide Leonard's choice of strategy since to move from the latter strategy profile to the former can only be achieved by a change in strategy choice by Jimmie. Although Leonard might strictly prefer the former to the latter, which of these two strategy profiles might be played should Leonard choose the strategy f_L is a choice that would be made by Jimmie and so presumably would be guided by his preferences, *not* Leonard's. We simply note at this juncture that the assumptions we impose on the preferences of players so as to establish the existence of an equilibrium of the ordinal game do not require any player to be able to express a preference between any pair of strategy profiles that involves a different strategy choice by any player other than himself.

3. Savage games

Following the discussion of the previous section, we now present the formal description of a Savage game. 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, which we take to be a compact metric space. 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 ith coordinate of z and z_{-i} is the projection of z into Z_{-i}.

Player i has a nonempty set F_i of A-valued functions on the state space Ω called the *strategy space*. A function $f_i : \Omega \to A$ 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 binary relation $\succsim_i$ on the set of strategy profiles F describing her weak preferences. As usual, $\succ_i$ and $\sim_i$ denote the asymmetric and symmetric parts of $\succsim_i$, respectively.

Our first assumption, A1, requires the preferences of the players to be complete and transitive *only* with respect to each player's own choices for a given (fixed) strategy profile of the other players.

A1. *For any $f \in F$ and any $g_i, h_i \in F_i$,*

(i) either $(f_i, f_{-i}) \succsim_i (g_i, f_{-i})$ or $(g_i, f_{-i}) \succsim_i (f_i, f_{-i})$, and

(ii) if $(f_i, f_{-i}) \succsim_i (g_i, f_{-i})$ *and* $(g_i, f_{-i}) \succsim_i (h_i, f_{-i})$, *then* $(f_i, f_{-i}) \succsim_i (h_i, f_{-i})$.

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 $i \in N$.

For any Savage game $(\Omega, A, (F_i, \succsim_i)_{i=1}^N)$, if for each i we consider a subset of strategies $\bar{F}_i \subseteq F_i$ and set $\bar{\succsim}_i$ to be the restriction of $\succsim_i$ to $\bar{F}$, then notice that, by construction, $(\Omega, A, (\bar{F}_i, \bar{\succsim}_i)_{i=1}^N)$ is itself a Savage game. By taking appropriate restrictions of the original game, we see that A1 has testable (revealed preference) implications for the equilibria of these restrictions.

Proposition 1. *If A1 holds, then for any* $f \in F$ *and any* $g_i \in F_i$, $f \succsim_i (g_i, f_{-i})$ *if and only if* f *is an equilibrium of the (restricted) game* $(\Omega, A, (\bar{F}_j, \bar{\succsim}_j)_{j=1}^N)$ *in which* $\bar{F}_i = \{f_i, g_i\}$ *and* $\bar{F}_{-i} = \{f_{-i}\}$.

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 $f_i, g_i : \Omega \to A$ 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.

One way to interpret the information structure of a player is to view the Savage game as a dynamic game. In the first stage, each player i receives partial information about the true state of the world, encoded in $\mathcal{F}_i$, the set of information events for player i. Only after that, in the *interim* stage, do players make their choice of action. Then in the last stage, all uncertainty is resolved and the final consequence is realized.

With this interpretation of the information structures of the players in mind, the next assumption can be viewed as being motivated by the standard consequentialist reasoning embodied in Savages's sure-thing principle. It is, in fact, based on the normative rule *dynamic programming solvability* introduced by Gul and Lantto (1990) in the context of individual choice under risk (that is, with exogenously specified probabilities).[3] Gul and Lantto highlight that although dynamic programming solvability constitutes a weakening of Savages's postulate P2, it still allows for a simplification of the task faced by an individual in finding an optimal "plan of action" for a two-stage decision tree by

[3] Grant et al. (2000) show how this property can be translated to the Savage framework.

allowing the player to "fold back" or "roll back" the two-stage decision tree. To illustrate the idea, Gul and Lantto give the example of a commuter having to decide how to go to work. The options are to walk, to drive, to bike, or to take the bus. Suppose the following two plans are optimal: (i) drive if it rains, bike if it is sunny, and (ii) take the bus if it rains, walk if it is sunny. Then, they argue, the following plans of actions should also be optimal: (iii) drive if it rains, walk if it is sunny, and (iv) take the bus if it rains, bike if it is sunny. In the setting of a Savage game this translates into the requirement that piecewise combinations of best responses (consistent with that player's information) remain a best response.

In what follows, 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 as an arbitrary member of F_i, and E always denotes an event in $\mathcal{F}_i$.

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

The reader will see in Section 5 that in games in which the preferences are given by payoffs, A2 holds when a wide range of assumptions that have been studied in the literature are satisfied. In particular and as has been already foreshadowed above, the following proposition establishes that A2 is implied by Savages's postulate P2.

Proposition 2. *The following condition implies A2:*

P2: If $(f_i, f_{-i}) \succsim_i (g_{iE} f_i, f_{-i})$, *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 nondegenerate in the sense of Savage. We make, however, the following "richness" assumption on strategies. It basically states that every state-contingent plan of action that a player can approximate by sequences of strategies must also be a feasible strategy for that player.

A3. *If* $E^n \in \mathcal{F}_i$ *is an increasing sequence of events for player i, then* $g_{i(\bigcup_n E^n)} f_i \in F_i$ *for every* $f_i, g_i \in F_i$.

Notice that, by construction, $\mathcal{F}_i$ is a collection of subsets of Ω (events) that contains Ω, and is closed under the operations of complement and finite unions, that is, $\mathcal{F}_i$ is an algebra of sets. The next proposition establishes that if assumption A3 holds, then $\mathcal{F}_i$ is also closed under countable unions, that is, it is a σ-algebra.

Proposition 3. *The collection of events for player i,* $\mathcal{F}_i$, *is an algebra over* Ω. *If A3 holds, then* $\mathcal{F}_i$ *is a* σ*-algebra.*

When the structure of a Bayesian game is available, we have the following corollary, which shows that $\mathcal{F}_i$ is precisely the smallest σ-algebra for which all player i's strategies, as usually defined, are measurable.

Corollary 1. *Let Σ_i be a σ-algebra over Ω and let 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. To do so, first, it is useful to define what it means for a pair of strategy profiles to be strategically equivalent for a player.

Strategic equivalence

The strategy profiles $f, g \in F$ are *strategically equivalent* for player i, denoted $f \approx_i g$, if for all $h_i, \hat{h}_i \in F_i$,

$$(h_i, f_{-i}) \succsim_i (\hat{h}_i, f_{-i}) \iff (h_i, g_{-i}) \succsim_i (\hat{h}_i, g_{-i}).$$

That is, if $f \approx_i g$, then player i's preferences over her own strategies in F_i are the same given the other players are choosing the profile f_{-i} as they are given the other players are choosing the profile g_{-i}.

The following properties readily follow from the definition of strategic equivalence.

Corollary 2. *The binary relation $\approx_i$ is an equivalence relation (that is, it is reflexive, symmetric, and transitive).*

Corollary 3. *For any $f \in F$ and any $g_i \in F_i$, $(f_i, f_{-i}) \approx_i (g_i, f_{-i})$ if and only if $(f_i, f_{-i}) \sim_i (g_i, f_{-i})$.*

We shall refer to the strategic analog of a null event as a strategically irrelevant event. An event will be deemed strategically irrelevant for a player if any deviation that player can make conditional on that event from any strategy profile leaves that player indifferent *and* does not affect any *other* player's preferences over her own strategies. That is, the original strategy profile and the strategy profile resulting from that player's deviation are *strategically equivalent for every player*.

Strategically irrelevant events

An event $E \in \mathcal{F}_i$ is *strategically irrelevant* for player i if, for all $f \in F$ and all $g_i \in F_i$, we have $(g_{iE} f_i, f_{-i}) \approx_j (f_i, f_{-i})$ for every player $j \in N$. Denote by $\mathcal{N}_i$ the set of all events that are irrelevant for player i. Let $\mathcal{R}_i = \mathcal{F}_i \setminus \mathcal{N}_i$ be the set of *strategically 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 strategically irrelevant 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 has any strategic relevance for any of the players. However, there exists at least one deviation by player j that is strategically relevant either for that player or for at least one of the other players. Even in the context of individual choice under uncertainty, Karni et al. (1983) note that if preferences are state-dependent, then interpreting null events as events 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.

Turning to continuity, we require the preferences of the players to be continuous with respect to statewise converging sequences of strategy profiles. Notice that condition A4 is satisfied in continuous Bayesian games in which the payoff of the players is computed by means of an integral.

A4. *If $f_i^n \in F_i$, $g_i^n \in F_i$, and $f_{-i}^n \in F_{-i}$ are sequences converging statewise to $f_i \in F_i$, $g_i \in F_i$, and $f_{-i} \in F_{-i}$, respectively, and $(g_i^n, f_{-i}^n) \succsim_i (f_i^n, f_{-i}^n)$ for all n, then $(g_i, f_{-i}) \succsim_i (f_i, f_{-i})$.*

By construction, the set of strategically irrelevant events for player i is closed with respect to subsets and finite unions, that is, it is an ideal. The next proposition establishes that, if assumption A4 holds, then $\mathcal{N}_i$ is also closed under countable unions, that is, it is a σ-ideal.

Proposition 4. *The collection of irrelevant events for player i, $\mathcal{N}_i$, 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 strategically relevant events. 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 the following statements:*

(i) If $E \in \bigcap_m \mathcal{S}_i^m$, then E is strategically irrelevant for player i.

(ii) If E^n is a sequence of strategically relevant events for player i and

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

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

Assumption A5 can be interpreted as follows. Condition (i) indicates that the sequence $\mathcal{S}_i^m$ comprises families of small events forming a neighborhood base for the subfamily of irrelevant events. The expression $(1/n)|\{1 \le k \le n : \omega \in E^k\}|$ is the average incidence of state ω arising from the sequence of events E^n. Keeping that in mind, condition (ii) then implies that players understand irrelevant events to be limits of decreasing sequences of relevant events.

Assumption A5 adapts a condition of Ryll-Nardzewski and Kelley that is sufficient for the existence of a measure on a σ-algebra (see the addendum to Kelley 1959). An alternative to assumption A5 would be to follow the approach taken by Arrow (1971), Fishburn (1970), and Villegas (1964), and construct a probability measure over a given σ-algebra based on a qualitative probability.[4] The reason we do not follow that approach is that it entails monotonicity assumptions, which in turn imply state-independent preferences. In the context of a Bayesian game, that implies that each player's ex post payoff cannot depend on the players' types, thus ruling out an important class of Bayesian games.

The following proposition establishes the equivalence between A5 and existence of a measure on player i's collection of events $\mathcal{F}_i$.

Proposition 5. *If A3 and A4 hold for all players, then the following statements are equivalent:*

(i) *Assumption A5 holds for player i.*

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

As an example, suppose $\Omega = [0, 1]$. Suppose also that for a particular player i the set $\mathcal{F}_i$ is the Lebesgue subsets of Ω, and $\mathcal{N}_i$ is the collection of subsets of zero Lebesgue measure, λ. In this case, one possibility is for the family of events $\mathcal{S}_i^m$ to be given by $\mathcal{S}_i^m = \{E \in \mathcal{F}_i : \lambda(E) \leq 1/m\}$. Then one example of a sequence of strategically relevant events satisfying A5(ii) is given by $\{E^{nm}\}$, in which $E^{nm} = [(m-1)/n, m/n)$, $n \geq 1$, and $m = 1, \dots, n$. In this example, the measure π_i whose existence is guaranteed by Proposition 5 could be *any* measure whose collection of zero measure sets coincides with the collection of zero Lebesgue measure sets, $\mathcal{N}_i$.

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

A6. *If $(f_i, f_{-i}) \not\approx_j (g_i, f_{-i})$ for some $j \in N$, then for each $h_i \in F_i$ there exist events $\{E^1, \dots, E^n\}$ such that $\bigcup_k E^k = \Omega$ and $(f_i, f_{-i}) \not\approx_j (h_{iE^k} g_i, f_{-i})$ for all k.*

As with P6, this assumption ensures that each relevant event can be split into two disjoint relevant events. Therefore, A6 implies that the measure π_i over $\mathcal{F}_i$ from Proposition 5 is atomless.

Proposition 6. *If A6 holds, then every $E \in \mathcal{R}_i$ contains two disjoint events in $\mathcal{R}_i$.*

The following proposition asserts that, in fact, under A3 and A4, assumptions A5 and A6 are equivalent to measure π_i being atomless.

[4] A binary relation $\geq$ is a qualitative probability if

(i) $\geq$ is a weak order (reflexive, complete and transitive)

(ii) $E \geq \varnothing$ for every event E, and

(iii) $E' \cap E = E'' \cap E = \varnothing$ implies $[E' \geq E'' \iff E' \cup E \geq E'' \cup E]$.

Proposition 7. *If A3 and A4 hold for all players, then the following statements are equivalent:*

(i) Assumptions A5 and A6 hold for player i.

(ii) The set $\mathcal{F}_i$ admits an atomless measure π_i such that $\pi_i(E) = 0$ implies $E \in \mathcal{N}_i$.

Notice that in condition (ii) some strategically irrelevant events of a player may be given a positive measure. The only requirement is that all events assigned zero measure are strategically irrelevant for that player. As we shall see in the proof, since we are dealing with countably additive measures, however, there exists a corresponding measure that is positive on strategically relevant events and zero on all strategically irrelevant events.

Our final assumption imposes a restriction on the best responses of players that is essential for existence of equilibrium in the Savage game. It plays the role of a compactness assumption in more standard equilibrium existence proofs. To be more specific, assumption A7 requires players to be able to find best responses that are not too erratic.

Regularity of best responses

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 statements:*

(1) For each $f \in F$, there is $g_i \in X_i$ satisfying $(g_i, f_{-i}) \succsim_i (f_i, f_{-i})$.

(2) Strategy X_i is countably distinguished and every sequence in X_i has a subsequence converging statewise to a strategy in X_i.

The significance of A7 can be seen in some well known examples of Bayesian games that fail to have pure-strategy equilibrium. The following example is taken from Radner and Rosenthal (1982).[5]

Example 1. There are two players, 1 and 2. Each player has two actions, thus $A_1 = A_2 = \{0, 1\}$. Payoffs are zero sum, according to the following payoff matrix:

		Player 2	
		0	1
Player 1	0	$(1, -1)$	$(-1, 1)$
	1	$(-1, 1)$	$(1, -1)$

[5]Another particularly striking example is presented in Khan et al. (1999). An argument similar to the one presented for Example 1 shows that A7 fails to hold.

The players' type space is the unit square $[0,1]^2$. The players' types are distributed uniformly on the triangle of the unit square given by $\{(t_1,t_2): 0 \le t_1 \le t_2 \le 1\}$. Radner and Rosenthal show that this game has no pure-strategy equilibrium. At any equilibrium, for each player i, conditional on his type being t, the probability of the other player choosing action 0 has to be $1/2$. That is, for a pair of pure strategies (f_1, f_2) to be an equilibrium in this game, it is necessary and sufficient that

$$\frac{\lambda([t_1,1] \cap f_2^{-1}(0))}{(1-t_1)} = \frac{\lambda([0,t_2] \cap f_1^{-1}(0))}{t_2} = \frac{1}{2}$$

for almost every t_1 and almost every t_2, with λ denoting the Lebesgue measure. Therefore, if it existed, at a pure-strategy equilibrium, the strategies played would be very erratic. In fact, it is possible to construct a sequence of strategies of player 2 to which player 1's (essentially unique) best responses become increasingly erratic, and thus cannot be contained in a subset of strategies of player 1 that satisfies A7. To illustrate that statement, we provide the first few elements of one possible such erratic sequence of strategy profiles. Consider the following sequence of strategies of player 2:

$$\begin{aligned}
f_2^1(t_2) &= 0 \\
f_2^2(t_2) &= \begin{cases} 0 & \text{if } t_2 \in [0, \frac{2}{3}) \\ 1 & \text{if } t_2 \in [\frac{2}{3}, 1] \end{cases} \\
f_2^3(t_2) &= \begin{cases} 0 & \text{if } t_2 \in [0, \frac{1}{3}) \cup (\frac{3}{4}, 1] \\ 1 & \text{if } t_2 \in [\frac{1}{3}, \frac{3}{4}] \end{cases} \\
&\cdots
\end{aligned}$$

The unique (up to a zero Lebesgue measure set) best responses of player 1 would then be

$$\begin{aligned}
f_1^1(t_1) &= 0 \\
f_1^2(t_1) &= \begin{cases} 0 & \text{if } t_1 \in [0, \frac{1}{3}) \\ 1 & \text{if } t_1 \in [\frac{1}{3}, 1] \end{cases} \\
f_1^3(t_1) &= \begin{cases} 0 & \text{if } t_1 \in [0, \frac{1}{6}) \cup (\frac{1}{2}, 1] \\ 1 & \text{if } t_1 \in [\frac{1}{6}, \frac{1}{2}] \end{cases} \\
&\cdots
\end{aligned}$$

It is possible to continue the sequence of strategy profiles $\{(f_1^1, f_2^1), (f_1^2, f_2^2), (f_1^3, f_2^3), \ldots\}$ such that each element (f_1^n, f_2^n) has exactly $n-1$ discontinuity points. Moreover, because each f_1^n is the (essentially) unique best response to player 2's strategy f_2^n, for A7 to be satisfied the entire sequence $\{f_1^n\}$ must be contained in X_1. However, this sequence has no convergent subsequence. ◊

Assumption A7(ii) 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) Strategy X_i is a finite set.

(b) Strategy 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.[6]

(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.[7]

(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.[8]

(f) Suppose that Ω is a measure space with σ-algebra $\mathcal{F}$, and that 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 statewise convergence, and there is a probability measure π on $\mathcal{F}$ 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 1. *If A1–A7 hold for all players, then an equilibrium exists.*

The proof of Theorem 1 is in the Appendix, but we conclude this section with a brief overview of the relevance of each assumption. We prove the theorem by showing that the product correspondence of the players' best responses has a fixed point. Assumptions A1 and A3–A5 guarantee that each player's best response correspondence is well defined and has a closed graph.[9] Assumptions A2 and A6 guarantee that the values of each player's best response correspondence is path connected, which plays the role of convexity in standard fixed point proofs.[10] Finally, assumption A7 is used to show that

[6] A family of functions F between metric spaces is called equicontinuous if for every $\varepsilon > 0$ there exists $\delta > 0$ such that $d_1(x, y) < \delta$ implies $d_2(f(x), f(y)) < \varepsilon$ for every $f \in F$. A function f between metric spaces is Lipschitz continuous if there exists a constant K, called Lipschitz constant, such that $d_2(f(x), f(y)) < Kd_1(x, y)$.

[7] A Radon measure μ is a measure on the σ-algebra of Borel sets of a Hausdorff topological space X that is both locally finite (every point of X has a neighborhood U for which $\mu(U)$ is finite) and inner regular (for any Borel subset B of X, $\mu(B)$ is the supremum of $\mu(K)$ over all compact subsets K of B). The weak* topology on a set $\mathcal{M}$ of measures on a measurable space $(\Omega, \mathcal{F})$ is the weakest topology such that all the linear functionals $L_f : \mu \mapsto \int_\Omega f \, \mathrm{d}\mu$ for $f : \Omega \to \mathbb{R}$ $\mathcal{F}$-measurable and bounded, are continuous.

[8] Let $(\Omega, \mathcal{F}, \mu)$ be a measure space, and let A be a Banach space. Then $L_\infty(\Omega, \mathcal{F}; \mu)$ is the set of (equivalence classes of) essentially bounded functions from Ω to A endowed with the essential supremum norm:

$$\|f\|_\infty = \inf\{c > 0 : |f(\omega)| \le c \text{ for } \mu\text{-a.e. } \omega\}.$$

[9] A correspondence $B : X \twoheadrightarrow Y$ between topological spaces is said to have a closed graph if the set $\{(x, y) \in X \times Y : y \in B(x)\}$ is a closed subset of $X \times Y$.

[10] A topological space X is path connected if any two given elements of it can be joined by a path, that is, given $x, y \in X$, there exists a continuous map $f : [0, 1] \to X$ such that $f(0) = x$ and $f(1) = y$.

the product correspondence of best responses has a subcorrespondence whose image is contained in a compact subset of the strategy space. Moreover, this subcorrespondence inherits the path-connectedness and closed-graph properties of the larger best response correspondence. Therefore, this allows us to apply the Eilenberg and Montgomery fixed point theorem to this subcorrespondence.

4. Matching with incomplete information

To illustrate the main concepts, definitions, and structures of the framework developed in the previous section, we present a fully specified example of a purely ordinal matching process that satisfies all of our assumptions, hence establishing the existence of an equilibrium match. This model can be viewed as an abstraction of the academic job market or the system of academic admissions. As far as we know, there are no other examples in the literature of incomplete information matching games without values, beliefs, and a specification of a utility representation for the preferences of the players.

There is a finite set of firms, $N = \{1, \dots, N\}$; we denote an individual firm by $i \in N$. There is also a finite set of workers, $W = \{1, \dots, W\}$; we denote an individual worker by $j \in W$. The end result of the matching process in this job market for worker j is denoted by a pair (i, b). The first element $i \in N \cup \{0\}$ is the firm to which this worker is allocated, with the interpretation that 0 corresponds to the worker not being employed by any of the firms in the market. The second element b specifies for each worker j, the remuneration bundle that firm i pays worker j.

To avoid dealing with measurability issues and for ease of exposition, we identify the space $\mathcal{B}$ of feasible remuneration bundles with a *finite* subset of some Euclidean space $\mathbb{R}^\ell$, with $\ell \geq 1$. As an example, one dimension could represent the salary offer toward which workers have strictly monotone preferences. Other dimensions could be the levels of health and retirement benefits. Still other dimensions could represent nonmonetary compensations, such as the city in which the job is to be performed. We assume $0 \in \mathcal{B}$ and take this to be the "default bundle" offered to any worker that is allocated to "firm" 0, as well as the default bundle offered by any firm to a worker with whom the firm does not want to be matched.

Hence formally, an outcome of this matching game can be taken to be a mapping

$$\theta : W \mapsto (N \times \mathcal{B}) \cup \{(0, 0)\},$$

where for each worker j, $\theta_1(j)$ is the firm at which worker j ends up employed and $\theta_2(j)$ is her remuneration bundle. Let Θ denote the (by construction, finite) set of all possible outcomes of the matching process.

So that workers need not consider any strategic interactions when making their choices, we assume that each worker only cares about the specifics of her own match. That is, worker j is indifferent between any pair of outcomes θ and θ' for which $\theta(j) = \theta'(j)$. In particular, each worker j in W is characterized by a binary relation P_j defined on $(N \times \mathcal{B}) \cup \{(0, 0)\}$ that is complete, transitive, and (so that best responses will be unique) asymmetric. That is, P_j is a (complete) strict preference relation. In addition, we suppose

$$(0,0)\ P_j\ (i,0) \quad \text{for every } i \in N \text{ and every } j \in W.$$

Asymmetry of P_j precludes $(i,0)\ P_j\ (0,0)$ from holding for any firm i and any worker j. It follows in the market game described below that an "offer" of the (default) bundle 0 by a firm to a worker guarantees that this worker will not be employed by that firm. The preference relation P_j of each worker j is taken to be common knowledge among the market participants.

The preferences of the firms, however, are not commonly known. The uncertainty they face at different stages of their market interactions requires us to define their preferences at three different levels of information. In reverse chronological order we have the ex post level, the interim level, and the ex ante level.

Firm i's private information is encoded in a (measurable) signal space (Ω_i, Σ_i), such that each Ω_i is a separable and complete metric space, and Σ_i is the Borel σ-algebra. Each (Ω_i, Σ_i) is assumed to be rich enough to admit an atomless measure. Let $\Omega = \times_i \Omega_i$ denote the product of the Ω_i's, and let Σ denote the (product) Borel σ-algebra.

Let $\mathcal{R}$ denote the (finite) set of complete and transitive binary relations defined on Θ. Given the profile of signals ω in Ω, each firm i's ex post preferences over outcomes Θ is given by $R_i^\omega \in \mathcal{R}$. We assume that the realization of the signal ω_i fully parametrizes (or reveals) the ex post preferences of firm i, that is, $R_i^\omega = R_i^{\omega'}$ whenever $\omega_i = \omega_i'$. We also assume that these preferences are measurable, that is, for each $R \in \mathcal{R}$, the set $T_i(R) = \{\omega_i \in \Omega_i : R = R_i^\omega, \text{for some } \omega_{-i} \in \Omega_{-i}\}$ is in Σ_i.

Let π be a probability measure on Σ. In this setting, this probability measure π need not be interpreted as a (common) prior belief; rather each of its marginals π_i over Σ_i simply identifies the sets of signal realizations for firm i that are strategically irrelevant. We further assume that each marginal π_i is atomless. Notice, however, that the assumptions we impose on this probability measure do not rule out the possibility of some correlation among the realization of preferences of the firms. At the interim level, on receipt of its signal realization ω_i, firm i has a conditional preference $R_i^{\omega_i}$ over the set Y_{-i} of Σ_{-i}-measurable functions from Ω_{-i} into Θ. Finally, at the beginning of the game, each firm i aggregates these conditional preferences in the form of an ex ante preference R_i over the set Y of Σ-measurable functions from Ω into Θ. Let $P_i^{\omega_i}$ (respectively, P_i) denote the strict preference relation derived from $R_i^{\omega_i}$ (respectively, R_i). We make two consistency assumptions on the preferences of the firms. The first is a consistency assumption between the ex post and the interim preferences. It essentially requires that if there are two signals of firm i, ω_i and ω_i', that reveal the same ex post preferences, then the corresponding interim preferences, $R_i^{\omega_i}$ and $R_i^{\omega_i'}$, cannot differ. The second assumption is a monotonicity requirement between the interim and the ex ante preferences. Formally, we assume that the preferences of the firms satisfy the following conditions:

(i) If $R_i^\omega = R_i^{\omega'}$, then $R_i^{\omega_i} = R_i^{\omega_i'}$.

(ii) For any $y, y' \in Y$, if $y(\omega_i, \cdot)\; R_i^{\omega_i}\; y'(\omega_i, \cdot)$ for π_i-almost every ω_i, then $y\; R_i\; y'$. Moreover, if additionally $y(\omega_i, \cdot)\; P_i^{\omega_i}\; y'(\omega_i, \cdot)$ for every $\omega_i \in T$ with $\pi_i(T) > 0$, then $y\; P_i\; y'$.

Finally, we assume that the firms' preferences are continuous with respect to sequences of strategy profiles that converge in measure. More specifically, we assume that the following continuity condition holds:

(iii) Fix a firm i, and let y^n and z^n be two sequences in Y that converge in π measure to y and z, respectively. If $y^n\; R_i\; z^n$ for every n, then $y\; R_i\; z$.

The matching game is played as follows. After each firm i observes ω_i, both sides of the market meet, and each firm i sends an offer of a remuneration bundle to each worker $j \in W$. Thus the action space of firm i is $\mathcal{B}^W$, where for each $a_i \in \mathcal{B}^W$, the bundle $a_i(j)$ is the one offered by firm i to worker j. For each worker j, given the action profile $a \in (\mathcal{B}^W)^N$, define her set of available actions as $A_j(a) = \{(0, 0)\} \cup \{(1, a_1(j)), \dots, (N, a_N(j))\}$. Worker j simply chooses the (unique) most preferred option from that set according to her strict preference relation P_j, regardless of what other workers are doing. The resulting outcome is the one corresponding to the accepted offers.

Unlike the interaction among workers, the interaction among the firms is strategic, and can be modeled as a Savage game described by $(\Omega, A, (F_i, \succsim_i)_{i\in N})$ as follows. In the game played by the firms, the common state space of the firms is Ω, the product space of signal realizations of the firms. The action space of each firm is $A = \mathcal{B}^W$. The set of strategies F_i available to firm i is the set of Σ-measurable functions from Ω to A, with the (informational) restriction that $f_i(\omega_i, \omega_{-i}) = f_i(\omega_i, \omega'_{-i})$ for every $\omega_{-i}, \omega'_{-i} \in \Omega_{-i}$.

Let $\phi : A^N \to \Theta$ denote the function that maps a given profile of offers from the firms to an outcome, determined by the decisions of the workers. That is,

$$\phi(a) = \theta \in \Theta,$$

such that for each worker j in W, $\theta(j) \in A_j(a)$ and

$$\theta(j)\; P_j\; \theta' \quad \text{for all } \theta' \in A_j(a) \setminus \{\theta(j)\}.$$

In this Savage game, given $f, g \in F = \times_{i=1}^N F_i$, firm i's preferences are given by

$$f \succsim_i g \quad \Longleftrightarrow \quad \phi \circ f\; R_i\; \phi \circ g.$$

To show that the Savage game played by the firms has an equilibrium it is enough to show it satisfies all of the assumptions A1–A7.

Clearly, assumption A1 is satisfied. To see that the assumption A2 holds, suppose that $(f_i, f_{-i}) \sim_i (g_i, f_{-i}) \succsim_i (h_i, f_{-i})$ for all $h_i \in F_i$. Consistency condition (ii) implies that $\phi \circ (g_i(\omega_i), f_{-i})\; R_i^{\omega_i}\; \phi \circ (f_i(\omega_i), f_{-i})$ for π_i-almost every ω_i. Thus for any event $E \in \Sigma_i$, $\phi \circ (g_{iE} f_i(\omega_i), f_{-i})\; R_i^{\omega_i}\; \phi \circ (f_i(\omega_i), f_{-i})$ for π_i-almost every ω_i. Applying again the consistency condition (ii) yields that $\phi \circ (g_{iE} f_i, f_{-i})\; R_i\; \phi \circ (f_i, f_{-i})$, thus

$(g_{iE}f_i, f_{-i}) \succsim_i (f_i, f_{-i})$, as required. By construction of the strategy sets F_i, assumption A3 also holds. The continuity assumption A4 is implied by the continuity requirement (iii). To see that assumption A5 is satisfied, take for each firm i the collection of events $\mathcal{S}_i^m = \{E \in \Sigma_i : \pi_i(E) \leq 1/m\}$.

To check A6, suppose that $(f_i, f_{-i}) \not\approx_j (g_i, f_{-i})$ and take any $h_i \in F_i$. For each $m = 1, 2, \dots$ take a partition $\mathcal{E}^m$ of Ω_i such that $\pi_i(E) = 1/2^m$ for every $E \in \mathcal{E}^m$. We claim that there exists an n such that $(f_i, f_{-i}) \not\approx_j (h_{iE}g_i, f_{-i})$ for every $E \in \mathcal{E}^n$. Suppose not, that is, suppose that for each m there exists $E^m \in \mathcal{E}^m$ such that $(f_i, f_{-i}) \approx_j (h_{iE^m}g_i, f_{-i})$. The sequence of strategy profiles $(h_{iE^m}g_i, f_{-i})$ converges in measure to (g_i, f_{-i}). By the continuity assumption (iii), it follows that $(f_i, f_{-i}) \approx_j (g_i, f_{-i})$, a contradiction. Finally, it is easy to see that A7 holds by defining for each firm i the set X_i to be the set of strategies that can be written as

$$f_i(\omega_i) = \sum_{R \in \mathcal{R}} a^R \chi_{T_i(R)}(\omega_i),$$

with $a^R \in A$. That is, X_i is the set of strategies of firm i that are *constant* on each subset $T_i(R)$ of Ω_i. Since $\mathcal{R}$, the set of ex post preferences defined on Θ, is a finite set, it follows that X_i is a *finitely* distinguished subset of the set of strategies of firm i. Moreover, since firm i's interim preferences are constant on each subset $T_i(R)$ of Ω_i, it follows that for each $f \in F$, there exists $g_i \in X_i$ satisfying $(g_i, f_{-i}) \succsim_i (f_i, f_{-i})$. Thus A7 is satisfied.

5. Sufficient conditions on recursive payoffs

The purpose of this section is to use the Savage game framework to study a class of games in which a player's ex ante evaluation of strategy profiles can be expressed as a function of her interim utility or payoff. To do this analysis, we introduce the concept of recursive payoffs, which is a decomposition of players' payoffs into interim utilities and ex ante utilities. Interim utilities capture the payoff a player of fixed type gets from her actions, given the strategy profile of other players. The ex ante utility of each player is calculated by means of an aggregator function, which captures only attributes of individual preferences, and is independent of the strategic environment. That is, interim utilities translate strategy profiles of other players into type-dependent utility payoffs, whereas the aggregator function evaluates these type-dependent utility payoffs according to each individual's risk preferences. The canonical example is Bayesian games, where it is always possible to express the players ex ante motivations in terms of own-type-wise maximization of an integral function over other players' types.

Because A7 is a game-specific assumption without an analog in the literature of individual decision-making under uncertainty, this section focuses on assumptions A1–A6. It is worth noting that it is easier in practice to establish assumptions A1–A6 independently of assumption A7.

Consider a game in *interim utility* form specified as

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

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 $\Sigma = \bigotimes_{i=1}^N \Sigma_i$.

Let F_i be the set of all Σ_i-measurable strategies $f_i : \Omega_i \to A_i$. Player i is associated with an *interim utility function* $\mathbf{V}_i : A_i \times F_{-i} \times \Omega_i \to \mathbb{R}$, 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 $\alpha_i : \Omega_i \to \mathbb{R}$ an *interim payoff* for player i. Player i has ex ante preferences over interim payoffs expressed by the utility function $\mathbf{W}_i$ that associates with each α_i an ex ante utility $\mathbf{W}_i(\alpha_i)$ in $\mathbb{R}$. The ex ante utility $\mathbf{U}_i(f)$ of player i for the strategy profile $f \in F$ is given by means of the recursive form

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

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

We make the following decomposition assumption on payoffs, which is essentially the translation of A2 to this setting.

B1. *For every $f_{-i} \in F_{-i}$ and $f_i, g_i \in F_i$, if $\mathbf{U}_i(f_i, f_{-i}) = \mathbf{U}_i(g_i, f_{-i}) \geq \mathbf{U}_i(h_i, f_{-i})$ for every $h_i \in F_i$, then $\mathbf{U}_i(f_{iE}g_i, f_{-i}) = \mathbf{U}_i(f_i, f_{-i})$ for every $E \in \mathcal{F}_i$.*

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

B2. *There exists a probability distribution $\mu : \Sigma \to [0, 1]$ such that the following statements hold:*

(i) All the marginal distributions $\mu_i : \Sigma_i \to [0, 1]$ of μ are atomless.

(ii) If f^n is a sequence of strategy profiles that converges μ-almost everywhere to f, then $\mathbf{U}_i(f^n)$ converges to $\mathbf{U}_i(f)$ 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 $\bar{F}_i = \{\bar{f}_i : f_i \in F_i\}$, 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 $\mathbf{U}_i(f) \geq \mathbf{U}_i(g)$ for any $\bar{f}, \bar{g} \in \bar{F}$.

Proposition 8. *If B1 and B2 hold, then the associated Savage game satisfies A1–A6.*

We now explore properties of players' ex ante attitudes toward interim payoffs as embodied in $\mathbf{W}_i$, which guarantee that B1 is satisfied, in the presence of B2. Of course,

$\mathbf{W}_i$ 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 nonstrategic uncertainty. In this regard, these properties can be compared to the generalizations of expected utility in the literature.

Example 2 (Dynamically consistent intertemporal payoffs). We say that the payoff $\mathbf{U}_i$ is dynamically consistent if

$$\mathbf{U}_i(f_i, f_{-i}) \geq \mathbf{U}_i(g_i, f_{-i}) \quad \text{for every } g_i \in F_i$$

implies that

$$\mathbf{V}_i(f_i(\omega_i), f_{-i}|\omega_i) \geq \mathbf{V}_i(g_i(\omega_i), f_{-i}|\omega_i) \quad \text{for } \mu_i\text{-almost every } \omega_i \text{ and every } g_i \in F_i.$$

That is, after player i receives partial information about the true state of the world, she will not want to revise a choice she made based on her ex ante payoff. We show that dynamically consistent payoffs satisfy B1.

For a given $f_{-i} \in F_{-i}$, take any $f_i, g_i \in F_i$ such that

$$\mathbf{U}_i(f_i, f_{-i}) = \mathbf{U}_i(g_i, f_{-i}) \geq \mathbf{U}_i(h_i, f_{-i}) \quad \text{for every } h_i \in F_i.$$

Dynamic consistency implies that

$$\mathbf{V}_i(f_i(\omega_i), f_{-i}|\omega_i) = \mathbf{V}_i(g_i(\omega_i), f_{-i}|\omega_i) \quad \text{for } \mu_i\text{-almost every } \omega_i.$$

Let $E = \{\omega_i : \mathbf{V}_i(f_i(\omega_i), f_{-i}|\omega_i) \neq \mathbf{V}_i(g_i(\omega_i), f_{-i}|\omega_i)\}$ and take any $E' \in \Sigma_i$. Then

$$\mathbf{U}_i(f_{iE'}g_i, f_{-i}) = \mathbf{W}_i(\mathbf{V}_i(f_{iE'}g_i, f_{-i})) = \mathbf{W}_i(\mathbf{V}_i(f_{i(E'\cap E)}g_i, f_{-i})) = \mathbf{U}_i(f_{i(E'\cap E)}g_i, f_{-i}).$$

Since $\mu_i(E) = 0$, assumption B2(ii) implies that

$$\mathbf{U}_i(f_{i(E'\cap E)}g_i, f_{-i}) = \mathbf{U}_i(g_i, f_{-i}),$$

which implies the required result. ◊

Example 3 (Strictly monotone utility). Some form of monotonicity is present in nearly all generalizations of expected utility. Using the measure from B2, if $\mathbf{W}_i$ is strictly monotone for the marginal μ_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 μ_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 μ_i measure. Suppose that $\mathbf{W}_i$ is strictly monotone, that is, $\mathbf{W}_i(\alpha_i) \geq \mathbf{W}_i(\beta_i)$ holds whenever $\alpha_i \geq \beta_i$ and $\mathbf{W}_i(\alpha_i) > \mathbf{W}_i(\beta_i)$ holds whenever $\alpha_i > \beta_i$. We show that B1 holds.

Take $f_i, g_i \in F_i$ such that $\mathbf{U}_i(f_i, f_{-i}) = \mathbf{U}_i(g_i, f_{-i}) \geq \mathbf{U}_i(h_i, f_{-i})$ for every $h_i \in F_i$, and $E \in \Sigma_i$. For interim payoffs α_i and β_i denote by $\alpha_i \vee \beta_i$ and $\alpha_i \wedge \beta_i$ the statewise supremum and infimum payoffs, and notice that strict monotonicity implies that $\mathbf{V}_i(f_i, f_{-i}) \wedge \mathbf{V}_i(g_i, f_{-i}) = \mathbf{V}_i(f)$ μ_i-almost everywhere. (Otherwise, for $E' = \{\omega_i : \mathbf{V}_i(f_i(\omega_i), f_{-i}|\omega_i) >$

$\mathbf{V}_i(g_i(\omega_i), f_{-i}|\omega_i)\}$, $\mathbf{V}_i(f_{iE'}g_i, f_{-i}) = \mathbf{V}_i(f_i, f_{-i}) \vee \mathbf{V}_i(g_i, f_{-i}) > \mathbf{V}_i(f) \wedge \mathbf{V}_i(g_i, f_{-i})$, which implies that $\mathbf{U}_i(f_{iE'}g_i, f_{-i}) > \mathbf{U}_i(f) \wedge \mathbf{U}_i(g_i, f_{-i})$.) Therefore,

$$\begin{aligned}\mathbf{U}_i(f) \geq \mathbf{U}_i(f_{iE}g_i, f_{-i}) &= \mathbf{W}_i(\mathbf{V}_i(f_{iE}g_i, f_{-i})) \geq \mathbf{W}_i(\mathbf{V}_i(f_i, f_{-i}) \wedge \mathbf{V}_i(g_i, f_{-i})) \\ &= \mathbf{W}_i(\mathbf{V}_i(f)) = \mathbf{U}_i(f),\end{aligned}$$

which implies the required result. ◊

Example 4 (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 $\mathbf{W}_i$ can be represented by a Choquet integral and exhibits Schmeidler's (1989) notion of uncertainty aversion, then $\mathbf{W}_i$ 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 statewise supremum and infimum payoffs. Assume that $\mathbf{W}_i$ is nondecreasing in the sense that if $\alpha_i(\omega_i) \geq \beta_i(\omega_i)$ for all ω_i, then $\mathbf{W}_i(\alpha_i) \geq \mathbf{W}_i(\beta_i)$. Now suppose that $\mathbf{W}_i$ satisfies supermodularity:

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

We show that B1 holds.

Take $f_i, g_i \in F_i$ such that $\mathbf{U}_i(f_i, f_{-i}) = \mathbf{U}_i(g_i, f_{-i}) \geq \mathbf{U}_i(h_i, f_{-i})$ for every $h_i \in F_i$, and $E \in \Sigma_i$. Notice that supermodularity implies that $\mathbf{W}_i(\mathbf{V}_i(f_i, f_{-i}) \wedge \mathbf{V}_i(g_i, f_{-i})) = \mathbf{W}_i(\mathbf{V}_i(f))$. Therefore,

$$\begin{aligned}\mathbf{U}_i(f) \geq \mathbf{U}_i(f_{iE}g_i, f_{-i}) &= \mathbf{W}_i(\mathbf{V}_i(f_{iE}g_i, f_{-i})) \geq \mathbf{W}_i(\mathbf{V}_i(f_i, f_{-i}) \wedge \mathbf{V}_i(g_i, f_{-i})) \\ &= \mathbf{W}_i(\mathbf{V}_i(f)) = \mathbf{U}_i(f),\end{aligned}$$

as required. ◊

Example 5 (Decomposable choice). Moving away from explicit monotonicity, we give a "betweenness" condition on preferences over interim payoffs that generalizes Savages's P2 postulate and that also satisfies the property B1 above. Consider the *decomposable choice* property of Grant et al. (2000), which in this setting translates to the following condition:

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

To see that this condition satisfies B1, take $f_i, g_i \in F_i$ such that $\mathbf{U}_i(f_i, f_{-i}) = \mathbf{U}_i(g_i, f_{-i}) \geq \mathbf{U}_i(h_i, f_{-i})$ for every $h_i \in F_i$. Notice that if $\mathbf{W}_i(\mathbf{V}_i(f)) = \mathbf{W}_i(\mathbf{V}_i(g_i, f_{-i})) \geq \mathbf{W}_i(\mathbf{V}_i(f_{iE}g_i, f_{-i}))$ for all $E \in \Sigma_i$, then it cannot be the case that $\mathbf{W}_i(\mathbf{V}_i(f)) > \mathbf{W}_i(\mathbf{V}_i(g_{iE}f_i, f_{-i}))$ for $E \in \Sigma_i$. ◊

Example 6 (Recursive payoffs with multiple priors). Suppose that player i has a bounded measurable *payoff function*

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

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 $\pi_i(\cdot \mid \cdot) : \Sigma_{-i} \times \Omega_i \to [0, 1]$, 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.[11] Player i's 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):

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

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

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

We make three assumptions.

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

C2. *There is a probability measure $\mu : \Sigma \to [0, 1]$ such that the following statements hold:*

(1) Each $\pi \in \bigcup_i D_i$ 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.*

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

Epstein and Marinacci (2007) characterize this last condition for the maxmin expected utility form in terms of a condition of Kreps (1979).

Proposition 9. *If the game with multiple priors satisfies C1, C2, and C3, then the associated game in interim utility form satisfies B1 and B2.* $\diamondsuit$

[11] 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 an equilibrium result in this section does not require the decomposability of priors into marginals and conditionals.

6. 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 recursive payoffs with multiple priors as in Example 6.

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 $\nu_i : \Omega \to \mathbb{R}_+$, 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 $u_i : A \times \Omega \to \mathbb{R}$ of player i is given by

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

with Lipschitz-continuous functions $P_i : A_i \times \Omega_i \to \mathbb{S}^n$, $Q_i : \Omega_i \to \mathbb{S}^n$, and $R_i : A_{-i} \times \Omega \to \mathbb{B}^{n+1}$, and $0 \le \gamma_i < \frac{1}{2}$. 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 interpret $P_i(a_i, \omega_i)$ 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 "antisocial" for other types, for example, $P_i(a_i, \omega_i) = a_i$. We assume that the inverse correspondence $P_i^{-1} : A_i \times \Omega_i \to A_i$ defined by

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

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

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

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 8 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

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

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

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

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 nonempty-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}. Let 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 $u_i : A \times \Omega \to \mathbb{R}$ 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 $M_i : A_{-i} \times \Omega \to \mathbb{S}^n$ 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 1. 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 $\mu_{-i} : \Omega_{-i} \times \Omega_i \to \mathbb{R}_+$ and $\nu_{-i} : \Omega_{-i} \times \Omega_i \to \mathbb{R}_+$ 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:

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

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, and C3. By Proposition 9 we need only show that A7 is satisfied.

Since D_i satisfies the rectangularity property it follows that 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

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

if and only if

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

But there is always a point in $\mathbb{S}^n$ satisfying $\|a_i - m_i(f_{-i}|\omega_i)\| \geq \|m_i(f_{-i}|\omega_i)\|$. Therefore, the value $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ of such a point a_i is 1. But the maximum of $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ is also 1. From this we conclude that a_i maximizes $\mathbf{V}_i(a_i, f_{-i}|\omega_i)$ if and only if $\|a_i - m_i(f_{-i}|\omega_i)\| \geq \|m_i(f_{-i}|\omega_i)\|$. In particular, the maximizers of $\mathbf{V}_i$ have the form

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

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

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

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.

7. Concluding remarks and related work

We conclude with a discussion of some related models as well as issues arising from the results we have derived in the framework of Savage games.

Strategic interaction between non-expected utility maximizers

The literature on modeling strategic interactions between individuals whose preferences may not conform to expected utility theory has mostly focused on normal-form games of complete information. Hence the only uncertainty the players face is with regard to the strategy choices of their opponents. Azrieli and Teper (2011) provide a very nice summary of this literature, comparing and contrasting the different approaches that have been taken.[12]

[12]Rather than reproduce their discussion here we refer the interested reader to Section 1.2 on p. 311 of their paper.

In their paper, however, Azrieli and Teper consider a class of games with incomplete information where the payoff of a player can depend on the state of nature as well as the profile of actions chosen by the players.[13] Each player can choose any state-contingent *randomization* over her set of available actions that is measurable with respect to her information, which in turn is characterized by a partition of a *finite* state space. They allow for players whose ex ante preferences over strategy profiles are generated by (fairly) arbitrary functionals defined on the state-contingent payoffs associated with each strategy profile. Assuming these functionals satisfy standard continuity and monotonicity properties, their main result is that an ex ante equilibrium exists in every game if and only if these functionals are quasi-concave.

Enriching the strategy space so that players can choose (consistent with their respective information) state-contingent randomizations over their actions allows Azrieli and Teper to establish the existence of an ex ante equilibrium in an incomplete information game with a finite state space, just as Anscombe and Aumann (1963) were able to characterize the class of preferences that admit a subjective expected utility representation in a setting with a finite state space. We note, however, that this is achieved at the cost of assuming that players have access to objective randomizing devices, a significant departure from the Savage approach in which all uncertainty is subjective.

Azrieli and Teper also note that the quasi-concavity of the functionals that generate the players' ex ante preferences over strategy profiles can readily be related to the property of ambiguity aversion in the Anscombe and Aumann two-stage setting where the subjective uncertainty is resolved first followed by an objective randomization over the final consequences. However, as Eichberger and Kelsey (1996) argue, there is no natural counterpart interpretation in Savages's setting of purely subjective uncertainty. Hence we do not find it surprising that for any Savage game with which there is associated a game in interim utility form, properties B1 and B2 neither imply nor require the quasi-concavity of the ex ante utility $\mathbf{W}_i$ from Section 5 above.

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 nonneutral attitudes toward ambiguity. However, Epstein and Wang's setting does not allow for interdependent preferences. Bergemann et al. (2014) 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.

[13]Kajii and Ui (2005) also model a class of games with incomplete information for the particular class of preferences that admit a multiple prior representation.

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, and Tan and da Costa Werlang (1988) and Börgers (1993). However, in many of these papers, rationality is expressed in terms of "state-independent expected utility."[14]

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, 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 et al. (1950), Radner and Rosenthal (1982), and Milgrom and Weber (1985), and related literature.[15] 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.

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.

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

[15]Other results include Balder (1988), Khan and Sun (1995), Khan et al. (2006), Loeb and Sun (2006, 2009), Fu et al. (2007), Fu (2008), Podczeck (2009), Khan and Rath (2009), Wang and Zhang (2012), and Greinecker and Podczeck (2013). Using the same techniques, it is possible to show the existence of pure equilibrium when there is a continuum of agents; see Schmeidler (1973) and Mas-Colell (1984).

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 that 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.[16] 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.[17]

Adapting the notion in Section 5, a game with a separable structure is an N-player game given by the 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 $\Sigma = \bigotimes_{i=1}^N \Sigma_i$. An *act* for player i is a Σ-measurable function $y: \Omega \to O_i$. 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* $\zeta_i : A \times \Omega \to O_i$, 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 $f_i : \Omega_i \to A_i$. 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

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

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, $\bar{F}_i = \{\bar{f}_i : f_i \in F_i\}$ for each i, and $\bar{f} \succsim_i^* \bar{g}$ if and only if $f \succsim_i^* g$.

[16] Aumann and Dreze (2009) develop a related idea in which subjective risk in a game uses available strategies. See also Section 8 of Hammond (2004).

[17] 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.

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

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

represents $\succsim_i$ over acts for every i. This is the case for example when $\succsim_i$ satisfies all of Savages's postulates and the "monotone continuity" assumption (Arrow 1971 and Villegas 1964) 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: Proofs

Proof of Proposition 2

Suppose that $f \sim_i (g_i, f_{-i}) \succsim_i (h_i, f_{-i})$ for every $h_i \in F_i$. In particular, $f \sim_i (g_i, f_{-i}) \succsim_i (g_{iE} f_i, f_{-i})$ 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 3

It is immediate that $\mathcal{F}_i$ contains $\varnothing$ 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 1

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 $E = \{\omega : g_{iE} f_i(\omega) \neq f_i(\omega)\}$, which is in $\sigma(F_i)$. Thus, $\Sigma_i = \mathcal{F}_i$ and A3 holds.

Proof of Proposition 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

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

which is a contradiction, because E is strategically irrelevant for player i.

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

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

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

Proof of Proposition 5

We can assume without loss of generality that each $\mathcal{S}_i^m$ also has 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

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

Proposition 10. *The following statements hold true:*

(i) If $E \in \mathcal{R}_i^m$, then $E \notin \mathcal{S}_i^m$.

(ii) We have $\bigcup_m \mathcal{R}_i^m = \mathcal{R}_i$.

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

Turning to (ii), suppose that $E \in \mathcal{R}_i$. Suppose by way of contradiction that $E \notin \mathcal{R}_i^m$ for all m. Event E is associated with $E^m \in \mathcal{S}_i^m$ such that $D^m = E \ominus E^m \in \mathcal{N}_i$. Let $D = \bigcup_m D^m$, 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 (iii) 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 11. *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 $E^1, E^2, \dots, E^n$ in $\mathcal{R}_i^m$ we have

$$\max_{\omega\in\Omega} \frac{1}{n} |\{1 \le k \le n : \omega \in E^k\}| > c.$$

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

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

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

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

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, $f_i : \Omega \to \mathbb{R}$ is in $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ if it is $\mathcal{F}_i$-measurable and bounded, and $g_i : \Omega \to \mathbb{R}$ is in the equivalence class $[f_i]$ if $\{\omega : f_i(\omega) \ne g_i(\omega)\}$ 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 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 $f_i \ge g_i$ if $\{\omega : g_i(\omega) > f_i(\omega)\}$ 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 12. *There exists $c > 0$ such that the for any f convex hull C^m in $L_\infty(\mathcal{F}_i|\mathcal{N}_i)$ of $\{\chi_E : E \in \mathcal{R}_i^m\}$ we have $\|f\|_\infty \ge c$. In particular, C^m is disjoint from $\mathbf{c}/2 - L_\infty^+(\mathcal{F}_i|\mathcal{N}_i)$.*

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 nonnegative 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 $E^n \in \mathcal{F}_i$, $\bigcup_n E^n = E$, 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 $E \in \mathcal{R} = \bigcup_m \mathcal{R}_i^m$. That is, π_i is the required measure.

Proof of Proposition 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 $(g_{iE} f_i, f_{-i}) \not\approx_j f$. Thus, there exists a sequence $\{E^1, \dots, E^k\} \subseteq \mathcal{F}_i$, $\bigcup_k E^k = \Omega$, satisfying $(f_{iE^k}(g_{iE} f_i), f_{-i}) \not\approx_j f$ 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 7

We need only show that (ii) implies (i). Let π_i be from condition (ii). 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 $\bigcap_m \mathcal{S}_i^m \subseteq \mathcal{N}_i$.

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

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

This shows that A5 holds.

We show that A6 also holds. Suppose that $(g_i, f_{-i}) \not\approx_j f$. Suppose by way of contradiction that for each m there is E^m satisfying $\pi_i(E^m) \le 1/m$ such that $(h_{iE^m}g_i, f_{-i}) \approx_j f$. 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, $(g_i, f_{-i}) \approx_j f$, which is a contradiction.

A fixed point theorem

We begin with a statement of a fixed point theorem and apply it to prove Theorem 1. 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 $B: F \twoheadrightarrow F$ 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 statements hold:

(i) 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 $h \in F$.

(ii) Domain F is sequentially closed in $L(S, T)$.

(iii) Mapping 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 2 (Corollary 2.3, Meneghel and Tourky 2013). *Let $B: F \twoheadrightarrow F$ be a decomposable μ-sequentially closed-graphed mapping. If for a compact and metrizable $X \subseteq F$ we have $X \cap B(f) \neq \varnothing$ for each $f \in F$, then B has a fixed point.*

Proof of Theorem 1

Assume that A7 holds.

For each i let $S_i = \Omega$ and $T_i = A$. Each $f \in F$ 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 7. We will assume that at least one player has a relevant event and that all nonzero π_i are probability measures. Let $\mathcal{F} = \bigotimes_{i=1}^N \mathcal{F}_i$ 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 $\mu : \mathcal{F} \to [0, 1]$ 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

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

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 $\widetilde{B} : \widetilde{X} \to \widetilde{X}$ 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 2 gives us the required equilibrium.

Proof of Proposition 8

Suppose that $f \in F$ and $g_i \in F_i$ satisfy

$$\mathbf{U}_i(f) = \mathbf{U}_i(g) \geq \mathbf{U}_i(h_i, f_{-i})$$

for all $h_i \in F_i$. By B1 this means that

$$\mathbf{U}_i(g_{iE} f_i, f_{-i}) = \mathbf{U}_i(f)$$

for every $E \in \Sigma_i$. Thus, A2 holds. That A3 is satisfied is a consequence of Corollary 1. Now $\mathbf{U}_i$ 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 1, $\bar{\Sigma}_i = \mathcal{F}_i$. Now the restriction of μ_i to $\bar{\Sigma}_i$ is atomless and if $\mu(E) = 0$, then E is strategically null for player i by (ii) of B2. By Proposition 7 assumptions A5 and A6 hold.

Proof of Proposition 9

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

Anscombe, Francis J. and Robert J. Aumann (1963), "A definition of subjective probability." *The Annals of Mathematical Statistics*, 34, 199–205. [669]

Arrow, Kenneth J. (1971), *Essays in the Theory of Risk-Bearing*. Markham, Chicago. [652, 672]

Aumann, Robert J. (1974), "Subjectivity and correlation in randomized strategies." *Journal of Mathematical Economics*, 1, 67–96. [671]

Aumann, Robert J. and Jacques H. Dreze (2008), "Rational expectations in games." *American Economic Review*, 98, 72–86. [644]

Aumann, Robert and Jacques Dreze (2009), "Assessing strategic risk." *American Economic Journal: Microeconomics*, 1, 1–16. [671]

Azrieli, Yaron and Roee Teper (2011), "Uncertainty aversion and equilibrium existence in games with incomplete information." *Games and Economic Behavior*, 73, 310–317. [643, 668]

Bade, Sophie (2011), "Ambiguous act equilibria." *Games and Economic Behavior*, 71, 246–260. [643]

Balder, Erik J. (1988), "Generalized equilibrium results for games with incomplete information." *Mathematics of Operations Research*, 13, 265–276. [670]

Bergemann, Dirk, Stephen Morris, and Satoru Takahashi (2014), "Interdependent preferences and strategic distinguishability." Princeton University Economic Theory Research Paper 008-2014. [669]

Börgers, Tilman (1993), "Pure strategy dominance." *Econometrica*, 61, 423–430. [670]

Brandenburger, Adam and Eddie Dekel (1987), "Rationalizability and correlated equilibria." *Econometrica*, 55, 1391–1402. [670]

Brandenburger, Adam and Eddie Dekel (1993), "Hierarchies of beliefs and common knowledge." *Journal of Economic Theory*, 59, 189–198. [669]

Debreu, Gerard (1960), "Topological methods in cardinal utility theory." In *Mathematical Methods in the Social Sciences* (Kenneth J. Arrow, Samuel Karlin, and Patrick Suppes, eds.), 16–26, Stanford University Press, Stanford. [671]

Denneberg, Dieter (1994), *Non-Additive Measure and Integral*. Kluwer Academic, Boston. [663]

Di Tillio, Alfredo (2008), "Subjective expected utility in games." *Theoretical Economics*, 3, 287–323. [669]

Dvoretzky, Aryeh, Abraham Wald, and Jacob 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 of America*, 36, 256–260. [670]

Eichberger, Jürgen and David Kelsey (1996), "Uncertainty aversion and preference for randomisation." *Journal of Economic Theory*, 71, 31–43. [669]

Epstein, Larry G. (1997), "Preference, rationalizability and equilibrium." *Journal of Economic Theory*, 73, 1–29. [670]

Epstein, Larry G. and Massimo Marinacci (2007), "Mutual absolute continuity of multiple priors." *Journal of Economic Theory*, 137, 716–720. [664]

Epstein, Larry G. and Martin Schneider (2003), "Recursive multiple-priors." *Journal of Economic Theory*, 113, 1–31. [664]

Epstein, Larry G. and Tan Wang (1996), "'Beliefs about beliefs' without probabilities." *Econometrica*, 64, 1343–1373. [669]

Fishburn, Peter C. (1970), *Utility Theory for Decision Making*. Wiley, New York. [653]

Fishburn, Peter C. (1982), "Multilinear utility on products of mixture sets." In *Foundations of Expected Utility*, 85–98, Reidel, Dordrecht, The Netherlands. [644]

Fu, Haifeng (2008), "Mixed-strategy equilibria and strong purification for games with private and public information." *Economic Theory*, 37, 521–532. [670]

Fu, Haifeng, Yeneng Sun, Nicholas C. Yannelis, and Zhixiang Zhang (2007), "Pure strategy equilibria in games with private and public information." *Journal of Mathematical Economics*, 43, 523–531. [670]

Gilboa, Itzhak (1987), "Expected utility with purely subjective non-additive probabilities." *Journal of Mathematical Economics*, 16, 65–88. [663]

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics*, 18, 141–153. [664, 667]

Gilboa, Itzhak and David Schmeidler (2003), "A derivation of expected utility maximization in the context of a game." *Games and Economic Behavior*, 44, 172–182. [644]

Grant, Simon, Atsushi Kajii, and Ben Polak (2000), "Decomposable choice under uncertainty." *Journal of Economic Theory*, 92, 169–197. [649, 663]

Greinecker, Michael and Konrad Podczeck (2013), "Spurious purification." Working paper. [670]

Gul, Faruk and Outi Lantto (1990), "Betweenness satisfying preferences and dynamic choice." *Journal of Economic Theory*, 52, 162–177. [649]

Hammond, Peter J. (2004), "Expected utility in non-cooperative game theory." In *Handbook of Utility Theory: Volume 2: Extensions*, 979–1063, Springer, New York. [671]

Harsanyi, John C. (1967–1968), "Games with incomplete information played by 'Bayesian' players. Parts I–III." *Management Science*, 14, 159–182, 320–334, 486–502. [641]

Kajii, Atsushi and Takashi Ui (2005), "Incomplete information games with multiple priors." *Japanese Economic Review*, 56, 332–351. [669]

Karni, Edi (1985), *Decision Making Under Uncertainty: The Case of State-Dependent Preferences*. Harvard University Press, Cambridge, Massachusetts. [671]

Karni, Edi, David Schmeidler, and Karl Vind (1983), "On state dependent preferences and subjective probabilities." *Econometrica*, 51, 1021–1031. [651]

Kelley, John L. (1959), "Measures on Boolean algebras." *Pacific Journal of Mathematics*, 9, 1165–1177. [652]

Khan, M. Ali and Kali P. Rath (2009), "On games with incomplete information and the Dvoretsky–Wald–Wolfowitz theorem with countable partitions." *Journal of Mathematical Economics*, 45, 830–837. [670]

Khan, M. Ali, Kali P. Rath, and Yeneng Sun (1999), "On a private information game without pure strategy equilibria." *Journal of Mathematical Economics*, 31, 341–359. [654]

Khan, M. Ali, Kali P. Rath, and Yeneng Sun (2006), "The Dvoretzky–Wald–Wolfowitz theorem and purification in atomless finite-action games." *International Journal of Game Theory*, 34, 91–104. [670]

Khan, M. Ali and Yeneng Sun (1995), "Pure strategies in games with private information." *Journal of Mathematical Economics*, 24, 633–653. [670]

Kreps, David M. (1979), "A representation theorem for 'preference for flexibility'." *Econometrica*, 47, 565–577. [664]

Kupka, Ivan (2005), "Continuous selections for Lipschitz multifunctions." *Acta Mathematica Universitatis Comenianae*, 74, 133–141. [666, 668]

Loeb, Peter and Yeneng Sun (2006), "Purification of measure-valued maps." *Illinois Journal of Mathematics*, 50, 747–762. [670]

Loeb, Peter and Yeneng Sun (2009), "Purification and saturation." *Proceedings of the American Mathematical Society*, 137, 2719–2724. [670]

Mas-Colell, Andreu (1984), "On a theorem of Schmeidler." *Journal of Mathematical Economics*, 13, 201–206. [670]

Meneghel, Idione and Rabee Tourky (2013), "A fixed point theorem for closed-graphed decomposable-valued correspondences." Preprint. Available at http://arxiv.org/abs/1306.3724. [676]

Mertens, Jean-François and Shmuel Zamir (1985), "Formulation of Bayesian analysis for games with incomplete information." *International Journal of Game Theory*, 14, 1–29. [669]

Milgrom, Paul R. and Robert J. Weber (1985), "Distributional strategies for games with incomplete information." *Mathematics of Operations Research*, 10, 619–632. [670]

Morris, Stephen and Satoru Takahashi (2012), "Games in preference form and preference rationalizability." Princeton University Economic Theory Research Paper 43-2012. [670]

Podczeck, Konrad (2009), "On purification of measure-valued maps." *Economic Theory*, 38, 399–418. [670]

Radner, Roy and Robert W. Rosenthal (1982), "Private information and pure-strategy equilibria." *Mathematics of Operations Research*, 7, 401–409. [654, 670]

Reny, Philip J. (2011), "On the existence of monotone pure strategy equilibria in Bayesian games." *Econometrica*, 79, 499–553. [643]

Riedel, Frank and Linda Sass (2014), "Ellsberg games." *Theory and Decision*, 76, 469–509. [643]

Savage, Leonard J. (1954), *The Foundations of Statistics*. Wiley, New York. [644, 645]

Schmeidler, David (1973), "Equilibrium points of nonatomic games." *Journal of Statistical Physics*, 7, 295–300. [670]

Schmeidler, David (1989), "Subjective probability and expected utility without additivity." *Econometrica*, 57, 571–587. [663]

Tan, Tommy C.-C. and Sérgio Ribeiro da Costa Werlang (1988), "The Bayesian foundations of solution concepts of games." *Journal of Economic Theory*, 45, 370–391. [670]

Villegas, Carlos (1964), "On qualitative probability σ-algebras." *The Annals of Mathematical Statistics*, 35, 1787–1796. [653, 672]

Wakker, Peter P. and Horst Zank (1999), "State dependent expected utility for Savage's state space." *Mathematics of Operations Research*, 24, 8–34. [671]

Wang, Jianwei and Yongchao Zhang (2012), "Purification, saturation and the exact law of large numbers." *Economic Theory*, 50, 527–545. [670]

Co-editor Johannes Hörner handled this manuscript.

Submitted 2015-1-15. Final version accepted 2015-6-3. Available online 2015-6-4.