Conditional, hierarchical, multi-agent preferences
(preliminary report)

PIERO LA MURA
STANFORD UNIVERSITY
GRADUATE SCHOOL OF BUSINESS
STANFORD, CA 94305
PLAMURA@LELAND.STANFORD.EDU

YOAV SHOHAM
STANFORD UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE
GATES BUILDING, ROOM 140
STANFORD, CA 94305
SHOHAM@CS.STANFORD.EDU

ABSTRACT. We develop a revealed-preference theory for multiple agents. Some features of our construction, which draws heavily on Jeffrey's utility theory and on formal constructions by Domotor and Fishburn, are as follows. First, our system enjoys the "small-worlds" property. Second, it represents hierarchical preferences. As a result our expected utility representation is reminscent of type constructions in game theory, except that our construction features higher order utilities as well as higher order probabilities. Finally, our construction includes the representation of conditional preferences, including counterfactual preferences.

1. INTRODUCTION

Two aspects of game theory are very evident nowadays. The first is that it has become an indispensable tool, not only in economics but in a variety of other disciplines as well, from philosophy and psychology to political science and computer science. The other is that game theory lacks comprehensive foundations of the scope and depth found in single-agent decision theory. Evidence of this limitation can be found in the many debates concerning backward induction problems, admissibility, and other examples in which the traditional game theoretic predictions are either paradoxical or ambiguous. It would be quite handy to have a Savage-style characterization of game theory, which would clarify the assumptions underlying different solution concepts, and therefore also the contexts in which each one applies.

What would such a theory look like? At a minimum it should involve a set of agents and a set of intuitive objects (such as events or acts), individual preferences over these objects for each agent that can plausibly be elicited from people via interrogation or observation, and a representation of each such individual preference ordering in terms of subjective probability and utility particular to the individual agent.

In this paper we will provide a theory that has these properties. Somewhat paradoxically, while our primary motivation is this multi-agent setting, the bulk of our construction can be explained already in the single-agent setting; the extension to the multi-agent setting then becomes obvious. Indeed, we are critical of existing foundations of decision theory (and in particular of Savage's framework), and believe that our theory provides better foundations. However, rather than waste our ammunition by attacking Savage's theory, whose shortcomings can be (and have been) well camouflaged, we simply note that we are not aware of any successful attempt to generalize Savage's framework to the multi-agent setting, and claim that it is no accident.

After such boasting we must introduce a caveat. One of the attractive features of Savage's framework is the treatment of causality, as embodied in the notion of an act. Although we believe that acts fit in quite naturally in the theory we are about to present and are working to incorporate them, in this version of the paper acts will play no role.

For this reason we are also not yet in a position to tackle the paradoxes of game theory with our framework (that too is next on our agenda).

The main ingredients of our construction can be relayed concisely by reference to several existing lines of research in decision theory, on which we draw liberally (the following will not make sense to the reader unfamiliar with the references, but the rest of the paper is self contained). From Jeffrey [Jef65] we borrow the scalable expected utility construction with the "small worlds" property. We then specialize and extend the framework. We first specialize it by constructing a particular algebra on which to define Jeffrey utilities, one that involves higher-order preferences (that is, preferences over preferences). Applying results due to Domotor [Dom78] we immediately get an expected-utility representation of higher-order preferences, albeit a problematic one. Among its chief deficiencies is the lack of account for the dynamics of preferences, or how preferences change in the face of new evidence (including counterfactual evidence - this turns out to be an important point). We then exploit the structure of our hierarchical construction, and, adapting and reinterpreting a relatively unknown construction due to Fishburn [Fis82], we strenghten the expect-utility representation and avoid these deficiencies. In both cases the resulting expected utility representation is in the spirit of existing type constructions in game theory [MZ85, HS96], but whereas these nest only probabilities our representation nests both probabilities and utilities. The work which comes closest to our approach is, to the best of our knowledge, [EW96]. Yet, our representation turns out to be quite different from the one in [EW96], and extends to conditional and counterfactual preferences.

The next section contains the bulk of the technical material, and presents our single-agent construction. In the following section we harvest the fruit of this construction by easily extending it to the multi-agent case. We conclude with a brief summary.

2. HIERARCHICAL CONDITIONAL PREFERENCES: THE SINGLE AGENT CASE

As we have said, most of the work in our construction is done already in the single-agent case, which we explain in this section. Before we begin the construction, let us point out three important ingredients in it:

1. "Small worlds" property: One need not be required to express preference only (or even at all) among entities that depend on objects which are sufficiently rich to resolve all ambiguity (by way of contrast, Savage's preferences are defined on acts, which in turn are defined relative to states which are such rich objects).

2. "Hierarchical preferences": One can assign preferences over preferences, as in preferring smoking to not smoking but wishing one didn't have that preference.

3. "Conditional preferences" (including counterfactual preferences): One should be able to specify how one's preferences change in the face of new evidence, including evidence given a prior probability of 0.

Although we believe these criteria to be desirable and are proud that our theory meets them, we don't request the reader to accept this desirability as self evident. We mention them now because it is helpful to keep them in mind when following the stages of the construction; in particular, the next three subsections correspond to these three criteria.

2.1. Foundation: Jeffrey utility. The technical development in this subsection is due to Jeffrey [Jef65] and Domotor [Dom78]. We start with a set of possible worlds W, and a finite Boolean algebra A of subsets of W. A preference ordering on A is a complete and transitive binary relation $\succsim$ between nonempty pairs $E, F \in A$.

Definition 1. *An expected utility representation of an ordering $\succsim$ on a finite algebra A is a pair (p, u), where $p : A \rightarrow [0,1]$ is a probability function and $u : A - \{\varnothing\} \rightarrow \mathbb{R}$ is a utility function such that:*

1. *For all nonempty $E, F \in A$, $u(E) \geq u(F)$ if and only if $E \succsim F$*
2. *$p(E) > 0$ for all nonempty $E \in A$, and $u(W) = 0$*
3. *$u(E)p(E) = \sum_k u(E_k)p(E_k)$ for any finite, measurable partition $\{E_k\}_{k=1}^K$ of E*

Notice the remarkable structure of this definition, as compared to the expected utility representations of von Neumann and Morgenstern, Savage, and others. Here preferences on events are represented by their relative utilities rather than their relative expected utilities, and the probabilities only serve to constrain the utilities via the last condition. One way to think about Jeffrey utilities is as simultaneously playing the role of traditional (e.g., Savage) utilities and traditional conditional expected utilities (or CEUs, where $CEU(E) = EU(E|E)$). This is a direct reflection of the "small worlds" property; since the events among which one expresses preference are under-specified states of the worlds, the utility of each event has an expectation flavor to it. The Jeffrey *value* of an event, defined as the product of its probability and utility, can be thought of as representing the event's standard (unconditional) expected utility. Note that in Jeffrey's framework probability and value are additive functions, but utility is not.

Remark 2. *Since the probability of any nonempty event is positive, we have that $p(F|E) = p(E \cap F)/p(E)$ is always well defined. Hence, (3) can be equivalently written in the conditional form $u(E) = \sum_k u(E_k)p(E_k|E)$.*

The question is whether there exist conditions on the ordering that guarantee the existence of an expected utility representation. Jeffrey [Jef65] gives a number of such conditions, which constrain not only the ordering but also the algebra itself. Similar conditions are provided by Bolker [Bol67]. These sets of conditions are sufficient to guarantee the existence of an expected utility representation, and have the additional advantage of being fairly intuitive. In the long version of this paper we present these conditions. Here, however, we proceed to present an alternative axiomatization, provided by Domotor [Dom78]. Domotor's axiomatization has two advantages - it constrains only the ordering but not the algebra, and is a necessary as well as sufficient condition for representability via Jeffrey utility (Theorem 4 below). However, it suffers the disadvantage of being highly technical and unintuitive, and so we introduce the relevant axiom by reference only.

Definition 3. *A preference ordering is* regular *if and only if it satisfies J2 in [Dom78].*

Theorem 4 [Domotor 1978]. *A preference ordering on a finite algebra is regular if and only if it admits a (real-valued) expected utility representation.*

Remark 5. *If the algebra is infinite, Theorem 4 continues to hold, the only difference being that probability and utility are allowed to take nonstandard (i.e., infinitesimal) values.*

2.2. A static expected utility representation of hierarchical preferences. We now proceed to construct a particular Jeffrey/Domotor structure, one that will capture the hierarchical nature of the agent's preferences while still enjoying the small-worlds property.

We first construct a large-worlds ontology, and then use it to define a hierarchical small-worlds framework. We start with a set W of possible worlds; a possible world is to be thought of as a rich object that completely captures the truth of all propositions, including the agent's preferences. Next, we introduce a function $\succsim$ which associates with each possible world a regular ordering over 2^W; $\succsim_w$ is to be thought of as extracting from each possible world w the agent's preferences at w.

Remark 6. *The reader familiar with modal logic will note the difference between this construction and Kripke-style possible-worlds semantics. From the conceptual point of view, in the latter a possible world settles on the truth value of objective facts, whereas here a possible world settles the truth value of all propositions, including the subjective ones. From the technical standpoint, in a Kripke structure each world is mapped by the accessibility relation to a set of worlds, whereas here each world is mapped by* $\succsim$ *to a total ordering on the power set of worlds.*

We are actually not interested in the entire algebra 2^W, but rather in a specific sub-algebra, A, which is defined as follows. We start with a finite Boolean algebra A_0 (a sub-algebra of 2^W, and ultimately also of A). A_0 is thought of as describing the objective events, ones that do not capture the agent's mental state. These are the objective events the agent is aware of, or is capable of imagining. There is no requirement that these "exhaust" the space of objective events in any sense.

Next, for all nonempty $E, F \in A_0$, let $[E \succ F] = \{w \mid E \succ_w F\}$, $[E \sim F] = \{w \mid E \sim_w F\}$, and $[E \succsim F] = [E \succ F] \cup [E \sim F]$.

Let B_0 be the set generated from propositions $[E \succ F]$, $[E \sim F]$, where $E, F \in A_0$, by closing off with respect to finite intersections. Then B_0 is a π-system, and contains both the empty set $\varnothing = [E \succ E]$ and $W = [E \sim E]$. Elements of B_0 correspond to partial ordering on A_0; B_0 describes the set of *zero-order preferences.*

Let $A_1 = A_0 \cup B_0$ be the algebra generated by propositions E where $E \in A_0$ or $E \in B_0$. Clearly, both A_0 and B_0 are sub-algebras of A_1. Again, let B_1 be the π-system of finite intersections of propositions $[E \succ F]$, $[E \sim F]$, where now $E, F \in A_1$. Elements of B_1 represent *first-order preferences.*

Now recursively define the n-th order $(n > 1)$ algebras and preferences as follows:

- $A_n = A_{n-1} \cup B_{n-1}$,
- B_n is the set of all finite intersections of propositions $[E \succ F]$, $[E \sim F]$, where $E, F \in A_n$.

Let $A = \cup_n A_n$ be the algebra generated by events $E \in A_n$, $n \geq 0$. Then $E \in A$ if and only if $E \in A_n$ for some n. Let B be the π-system generated by $[E \succ F]$, $[E \sim F]$, where $E, F \in A$. Notice that $B \subset A$, and hence $A = A \cup B$. Therefore, further iteration is superfluous: all the preferences on events in A are already included in A.

Remark 7. *Recall that A is the sub-algebra of 2^W in which we are interested. It might be asked why not define the preferences $\succsim_w$ only on A, and turn the above construction into a fixpoint definition. It turns out, however, that in a later development in this paper (when we define mixture operations) we will need to include events outside A.*

We are now close to achieving our first goal, an expected-utility representation of hierarchical preferences. What we are after is, for each w, giving the ordering $\succsim_w$ an expected utility representation. Our work is done almost automatically by the Jeffrey/Domotor result. We simply need to note the following:

Lemma 8. *If an ordering over an algebra is regular, so is its projection to any sub-algebra.*

Since A is a sub-algebra of 2^W, for every $w \in W$ the restriction of $\succsim_w$ to A is also regular, and therefore admits an expected utility representation.

So it would seem that we have accomplished our goal, but in fact there are several interrelated reasons for dissatisfaction:

- The requirement that every possible (i.e., nonempty) event be given a non-zero probability is conceptually problematic, since it doesn't allow the agent to recognize certain events as meaningful (or "possible") and disbelieve them at the same time. In particular, in the multi-agent setting, this will prohibit representing dominated strategies as actual but disbelieved possibilities.
- Beyond the conceptual difficulty, the above requirement has unpleasant technical ramifications. In particular, since A is in general infinite, the representation we have uses nonstandard (i.e., infinitesimal) probabilities and utilities.
- The current theory does not account for the way in which preferences (and thus probabilities and utilities) change in the face of new information; for this reason we term it "static." In particular, there is no obvious role for Bayesian conditioning, and no account of counterfactual conditioning.
- Perhaps most damningly, the current theory really does not make use of the hierarchical construction, beyond the weak use in the Lemma 8. In particular, nothing in the theory constrains the relationship between preferences at different levels, contradicting intuition that such "coherence" constraints ought to exist.

We now proceed to develop a theory that does not have these shortcomings.

2.3. A dynamic expected utility representation of hierarchical preferences. Recall that the expected utility representation afforded by the Jeffrey/Domotor construction involves probabilities and utilities that obey the following equation:

$$u(E) = \sum_k u(E_k)p(E_k|E) \tag{1}$$

First on our agenda is to strenghten this property, and ensure that the probability and utility obey the equation

$$u(E) = \sum_k u(E_k)p_E(E_k) \tag{2}$$

where p is a *conditional probability system (CPS)*. Recall that, given an algebra A, a CPS (aka Popper function) p assigns to every non-empty conditioning event $E \in A$ a probability function over $A \cap E$. Furthermore, p_E agrees with Bayesian conditioning whenever possible: for any nonempty $E, F, G \in A$, such that $G \subset F \subset E$, $p_E(G) =$

$p_E(F)p_F(G)$. If $E = W$, and the unconditional probability of F is positive (that is, $p(F) = p_W(F) > 0$), then the above formula yields

$$p_F(G) = p(G \mid F) = \frac{p(G \cap F)}{p(F)}.$$

If we manage to guarantee (2) we will have escaped the first two limitations of the (1)-based representation discussed above. Now let's go a step further. Consider any $E \in A_n$ and $F \in B_n$. Our claim is that, given the intuition behind our construction, E ought to be probabilistically independent of F; the lower order events do not determine the higher order preferences, and vice versa. From this it follows that our expected utility representation should validate the following property:

$$u(E \cap F) = \sum u(E_k \cap F)p_E(E_k), \tag{3}$$

Note that (1) is obtained as a special case of (3) by selecting $p(E) > 0$ and $F = W$. Note also that we have now escaped the third and fourth limitations of (1)-based representation.

How do we obtain (3)? We do so by leveraging a relatively unknown construction due to Fishburn [Fis82]. His motivation was different from ours – giving a conditional version of Savage's construction. However, we will adapt and reinterpret the mathematics to fit our intended interpretation. Fishburn starts with preferences defined on pairs (x, E), where the first argument is an act, and the second an environmental event.[1] In our interpretation, events $E \cap F$, where $E \in A_n$ and $F \in B_n$ play the role of (act,event) pairs (F, E). In other words, acts are viewed as being themselves events: they represent (generally incomplete) descriptions of the agent's conditional preferences (and hence beliefs), and characterize conditional revealed-preference behavior.

We proceed now with the technical construction.

Mixtures. Let R be the set of all expected utility representations (p, u) on some algebra A.

For any two representations $(p', u'), (p'', u'') \in R$, and for any $\lambda \in [0, 1]$, we define their $(\lambda-)$mixture to be a new representation $(p, u) = (p', u')\lambda(p'', u'')$ such that:

$$p(E) = p'\lambda p''(E) = \lambda p'(E) + (1 - \lambda)p''(E)$$

$$u(E) = u'\lambda u''(E) = \frac{\lambda u'(E)p'(E) + (1 - \lambda)u''(E)p''(E)}{p(E)}.$$

For any two nonempty subsets of representations $x, y \in 2^R$, we define their λ-mixture as the (nonempty) subset
$x\lambda y = \{(p, u) \in R \mid (p, u) = (p', u')\lambda(p'', u''),\ (p', u') \in x,\ (p'', u'') \in y\}.$

[1] We don't discuss Fishburn's inuition at length here, both because ours is different and because his is perhaps problematic. Briefly, however, the event is taken from an algebra on a set of states of nature, and the act is taken from a mixture set. An act x can be thought of as a probability measure on environmental events, and in this interpretation the representation is an interesting mix between von Neumann-Morgenstern and Savage: the agent chooses objective lotteries, and Nature chooses the context in which the lottery is performed. While the lotteries are objective, the probability of an (external) event in the context of another event is subjective, and can be uniquely derived from the agent's preferences on pairs (x, E). Unfortunately, some of the hypotheses introduced by Fishburn in order to derive his representation are quite unappealing in the suggested interpretation, and this is probably why the result never gained much popularity.

Definition 9. *A nonempty subset $x \in 2^R$ is (mixture) convex if, for any $\lambda \in [0,1]$, $x = x\lambda x$.*

Let $M \subset 2^R$ be the set of all convex subsets of R. M has the useful property of being closed with respect to mixtures, as the following proposition shows.

Proposition 10. *For any $x, y \in M$, and for any $\lambda \in [0,1]$, $x\lambda y \in M$.*

Many decision-theoretic treatments postulate that the set of objects on which preferences are defined is a *mixture set*:

Definition 11. *A set X is a mixture set with respect to a mixture operation $x\lambda y$ if, for any $x, y \in X$ and $\lambda, \mu \in [0,1]$,*

1. $x1y = x$
2. $x\lambda y = y(1-\lambda)x$
3. $(x\lambda y)\mu y = x(\lambda\mu)y$

The set M defined above turns out to have the desired structure.

Proposition 12. *M is a mixture set.*

Axiomatization. As before, let M be the set of all convex subsets of R. For any $x \in M$, let $[x] = \{w \mid \text{there exist } (p,u) \in x \text{ that represents } \succsim_w\}$.

Remark 13. *Note that the overloading of the $[\ldots]$ operator is quite helpful. To begin with, $[E \succ F]$ and $[x]$ are of the same type, but the relationship is even tighter. Define $|E \succ F| = \{(p,u) \mid u(E) > u(F)\}$ and $|E \sim F| = \{(p,u) \mid u(E) = u(F)\}$. Note that both these types of set are convex, as are their finite intersections. Moreover, since $E \underset{(\sim)}{\succ} F$ if and only if $u(E) \underset{(=)}{>} u(F)$, then $[E \underset{(\sim)}{\succ} F] = [|E \underset{(\sim)}{\succ} F|]$.*

With this machinery, given a regular preference ordering $\succ$ we can define the induced partial ordering $\succ^*$ on pairs (x, E), where $x \in M$ and $E \in A$.

- $(x, E) \succ^* (y, F)$ if and only if $[x] \cap E \succ [y] \cap F$
- $(x, E) \sim^* (y, F)$ if and only if $[x] \cap E \sim [y] \cap F$

Remark 14. *Notice that (x, E) is ranked if and only if $[x] \cap E$ is nonempty. Hence, it is immediately verified that $\succ^*$ is asymmetric and transitive, and $\sim^*$ is symmetric and transitive. Yet, $\succ^*$ is not negatively transitive, and $\sim^*$ is not the symmetric complement of $\succ^*$.*

We assume that mixtures of possible preferences are also possible.

A0. If $x, y \in M$ correspond to nonempty $[x], [y]$, then $[x\lambda y]$ is also nonempty for all $\lambda \in (0,1)$.

Next, we introduce four additional axioms. The first explicitly relates to mixtures.

$A1$. (Substitution) For all $E, F \in A$ and $x, y, z, t \in M$, if $(x, E) \sim^* (z, F)$ and $(y, E) \sim^* (t, F)$ then $(x\frac{1}{2}y, E) \sim^* (z\frac{1}{2}t, F)$.

The second axiom also relates to mixtures, and ensures a classical (i.e., standard) representation.

$A2$. (Archimedean) $\{\alpha : (x\alpha y, E) \succsim^* (z, F)\}$ and $\{\beta : (z, F) \succsim^* (x\beta y, E)\}$ are closed subsets of $[0, 1]$.

We remark that $A2$ is imposed in order to obtain a classical representation: its main role is to ensure that probabilities and utilities can be taken to be standard-valued.

The third axiom is fairly uninteresting. Its main role is to avoid triviality.

$A3$. (Relevance) There exist $x, y \in M$ such that $(x, W) \succ^* (y, W)$.

Finally, the fourth axiom is a consistency requirement.

$A4$. (Consistency) For all $E, F, G \in A$, $G \subset [E \succ F]$ implies $G \cap E \succ G \cap F$.

The axiom says, somewhat tautologically, that in the context of the agent's preference of E to F (i.e., whenever $G \subset [E \succ F]$), E is indeed preferred to F.

Theorem 15. *Under $J2$ and $A0 - A4$, there exists a conditional representation (p, u) such that:*

1. *$p : A \times (A - \{\varnothing\}) \to [0, 1]$ is a conditional probability system*
2. *u is a real-valued utility function, defined for all $(x, E) \in M \times A$ such that $[x] \cap E \neq \varnothing$, with the following properties:*
   - *$u(x, E) \underset{(=)}{>} u(y, F)$ if and only if $[x] \cap E \underset{(\sim)}{\succ} [y] \cap F$*
   - *$u(x, E) = \sum u(x, E_k) p_E(E_k)$ for any finite, measurable partition $\{E_k\}$ of E.*

Remark 16. *If F is an element of B, the representation speclializes to $u'(F \cap E) = \sum u'(F \cap E_k) p_E(E_k)$, where u' is defined by $u'([x] \cap E) = u(x, E)$, whenever $F \cap E_k \neq \varnothing$ for all k. Furthermore, if we take F to be the whole set W, we get $u'(E) = \sum u'(E_k) p_E(E_k)$.*

3. MULTI-AGENT CONSTRUCTION

The construction introduced in the previous section can be easily generalized to the multi-agent case. As was mentioned in the introduction, although this the multi-agent case is our primary motivation, this section is brief because it is a straightforward extension of the single-agent case.

Let $I = \{1, ..., n\}$ be a set of agents. Agent $i \in I$ is assumed to have preferences (and hence beliefs) not only about the basic events in a finite algebra A_0 and its own preferences, but on other agents' preferences as well. (In the treatment here we have all agents share the base algebra A_0, though this can be relaxed.)

As in the single-agent case, W represents a set of possible worlds, and $\succsim_w^i$ is a function associating to each pair (i, w) a regular preference ordering on 2^W.

For any $E, F \in A_0$, we denote by $[E \succ^i F]$ the proposition $\{w \in W \mid E \succ_w^i F\}$ ("i (strictly) prefers E to F"), and by $[E \sim^i F]$ the set $\{w \in W \mid E \sim_w^i F\}$ ("i is indifferent between E and F"). We denote by B_0^i the π-system obtained by taking all finite intersections of propositions $[E \succ^i F]$ and $[E \sim^i F]$.

We recursively define n-th order $(n > 1)$ algebras and preferences:

- $A_n = A_{n-1} \cup (\cup_{i \in I} B^i_{n-1})$,
- B^i_n $(i \in I)$ is the set of all finite intersections of propositions $[E \succ^i F]$, $[E \sim^i F]$, where $E, F \in A_n$.
- $A = \cup_n A_n$ is the algebra generated by events $E \in A_n$ $(n \geq 0)$, and B^i is the π-system generated by $[E \succ^i F]$, $[E \sim^i F]$, where $E, F \in A$.

Again, we obtain a Jeffrey-style representation of agent i's preferences, for all A_n and $i \in I$:

$$u^i(E)p^i(E) = \sum u^i(E_k)p^i(E_k).$$

Moreover, under the same conditions introduced in the single-agent case, we get a Fishburn-style representation for pairs $(E, F) \in A \times B^i$, which satisfies

$$u^i(E \cap F) = \sum u^i(E_k \cap F)p^i_E(E_k)$$

whenever $E \cap F \neq \varnothing$.

4. CONCLUSIONS

We presented a decision-theoretic approach aimed at overcoming several well-known limitations of existing constructions, limitations that become particularly apparent – and disturbing – in multi-agent applications. Our approach enjoys several advantageous features, including the ability to represent:

- Preferences on incomplete descriptions of the world.
- Conditional behavior, even contingent on disbelieved (counterfactual) events.
- Higher-order preferences (and hence beliefs).

The current limitations include:

- The axioms are not optimized for the proposed interpretation. That is, we glue together constraints drawn from Domotor (or Jeffrey) and from Fishburn. Together these are sufficient to guarantee the representation we seek, but there is no reason to believe that they are necessary.
- we do not account for 'agent causality', or actions.
- As a result, we are not yet in a position to apply our construction to game theoretic situations.

We view the first limitation as more of a mathematical annoyance than anything else, and are actively working on removing the second one. We leave the explicit application of our theory to game theoretic problems to a future paper.

REFERENCES

[Bol67] Bolker, E.D., 1967, "A Simultaneous Axiomatization of Utility and Subjective Probability", *Philosophy of Science* 34, 333-340.

[Dom78] Domotor, Z., 1978, "Axiomatization of Jeffrey Utilities", *Synthese* 39, 165-210.

[EW96] Epstein, L. and Wang, T., "'Beliefs about Beliefs' without Probabilities", *Econometrica* 64, n.4, 1343-74.

[Fis82] Fishburn, P.C., 1982, *The Foundations of Expected Utility*, Reidel, Dordrecht, Holland.

[Jef65] Jeffrey, R.C., 1965, *The Logic of Decision*, University of Chicago Press, Chicago.

[HS96] Heifetz, A. and Samet, D., 1996, "Topology-Free Typology of Beliefs", ewp-game/9609002.

[MZ85] Mertens, J.-F. and Zamir, Z., 1985, "Formulation of Bayesian Analysis for Games with Incomplete Information", *International Journal of Game Theory* 14, 1-29.

[Sav54] Savage, L.J., 1954, *The Foundations of Statistics*, Wiley, New York.

[vNM44] von Neumann, J. and Morgenstern, O., 1944, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, New Jersey.