The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website http://www.econometricsociety.org or in the back cover of Econometrica). This statement must the included on all copies of this Article that are made available electronically or in any other format.

Econometrica, Vol. 69, No. 4 (July, 2001), 891–934

REPRESENTING PREFERENCES WITH A UNIQUE SUBJECTIVE STATE SPACE

BY EDDIE DEKEL, BARTON L. LIPMAN, AND ALDO RUSTICHINI[1]

We extend Kreps' (1979) analysis of preference for flexibility, reinterpreted by Kreps (1992) as a model of unforeseen contingencies. We enrich the choice set, consequently obtaining uniqueness results that were not possible in Kreps' model. We consider several representations and allow the agent to prefer commitment in some contingencies. In the representations, the agent acts as if she had coherent beliefs about a set of possible future (ex post) preferences, each of which is an expected-utility preference. We show that this set of ex post preferences, called the *subjective state space*, is essentially unique given the restriction that all ex post preferences are expected-utility preferences and is minimal even without this restriction. Because the subjective state space is identified, the way ex post utilities are aggregated into an ex ante ranking is also essentially unique. Hence when a representation that is additive across states exists, the additivity is meaningful in the sense that all representations are intrinsically additive. Uniqueness enables us to show that the size of the subjective state space provides a measure of the agent's uncertainty about future contingencies and that the way the states are aggregated indicates whether these contingencies lead to a desire for flexibility or commitment.

KEYWORDS: Unforeseen contingencies, preference for flexibility.

1. INTRODUCTION

1.1. *A Brief Overview of the Results*

KREPS (1979) SHOWED THAT THE PREFERENCE of an agent over sets of possible future choices or actions[2] can be represented using *subjective states* that are interpreted as the agent's (implicit) view of future possibilities. More precisely, a subjective state is a possible ex post preference over actions that will govern the agent's choice tomorrow of an action from the set she chooses today. Surprisingly, he showed that weak axioms on preferences were sufficient to give a representation in which the agent has a coherent subjective state space without assuming an exogenously given state space. Kreps (1992) reinterpreted this as a model of unforeseen contingencies. After briefly reviewing our results, we discuss this interpretation, which is our main motivation for this work, and then return to a detailed description of Kreps' work and our results.

[1] We thank Jeff Ely, Matt Jackson, Peter Klibanoff, George Mailath, Jean-François Mertens, Sujoy Mukerji, Klaus Nehring, Phil Reny, Shuyoung Shi, Rani Spiegler, Jeroen Swinkels, numerous seminar audiences, and Drew Fudenberg and three anonymous referees for helpful comments. Dekel and Rustichini thank the NSF and Lipman thanks SSHRCC for financial support for this research. Lipman also thanks Northwestern and Carnegie Mellon for their hospitality while this work was in progress. This paper was previously titled "A Unique Subjective State Space for Unforeseen Contingencies." This work was begun while the second author was at the University of Western Ontario, and while the third author was at CORE and CentER.

[2] For example, a specification of control rights in a firm can be interpreted as such a set since it specifies a restriction on the future actions of an agent.

Unfortunately, the subjective state space Kreps derived is not pinned down by the preferences, making it problematic to interpret it as the agent's view of what is possible and leading to other difficulties discussed below. We enrich the choice set to consist of sets of lotteries over future actions. We show that the subjective state space is *unique* whenever we can represent the agent's ex ante choice of a set of lotteries under the hypothesis that her ex post choice from the selected set satisfies the expected-utility axioms. We also show that such a representation is possible given a surprisingly weak condition: the decision maker must be indifferent to having the extra option of randomizing over the lotteries in a chosen set. Normatively and descriptively, this seems like a weak requirement. Thus, under mild assumptions, one does not need to assume the existence of an exogenous state space to deduce that decision makers will behave as if they have a unique such state space in mind.

The uniqueness result enables us to show that the agent's "uncertainty" about the future can be measured by the size of her subjective state space. It also enables us to identify the *aggregator*—that is, the way the agent aggregates her possible ex post utility levels into an ex ante evaluation. In particular, we also characterize when the representation is inherently *additive* across states.[3]

Without the restriction to ex post preferences that are EU, the subjective state space is not unique. However, we show that within a wide class of preferences, the EU subjective state space we derive is the smallest possible subjective state space.

In addition to allowing for lotteries, we modify Kreps' assumptions in another direction. Kreps derived the subjective state space by analyzing when flexibility was valued by the agent. Kreps assumed that flexibility was never disadvantageous—that is, that preferences are monotonic in the sense that a larger set is always weakly better. However, the agent might have in mind some situations where flexibility is costly. For example, the agent may envision a scenario in which commitment is valuable for strategic reasons. Alternatively, as recently proposed by Gul and Pesendorfer (1999), also in a sets-of-lotteries framework, the agent may derive disutility from temptations that might arise. Finally, the agent may simply find larger sets more difficult to analyze because of complexity considerations. Motivated in part by the contribution of Gul and Pesendorfer, we drop the monotonicity assumption for most of our analysis. Unlike them, we do not specify the particular form of the violations of monotonicity allowed. It turns out that the representation and uniqueness results do not require restricting the agent to only conceive of circumstances in which flexibility is valuable. In particular, our identification of the aggregator implies that we also uniquely identify those ex post contingencies in which the agent prefers flexibility and those in which she prefers commitment.

[3] A representation with a particular subjective state space is inherently additive if it must be a monotone transformation of an additive representation using (essentially) that same state space.

1.2. *The Unforeseen Contingencies Interpretation*

We are interested in a model that allows for unforeseen contingencies, in the sense that the agent does not have an exogenously given list of all possible states of the world. This may be because she sees some relevant considerations, but knows there may be others that she cannot specify. For example, perhaps she sees that a particular variable x is relevant, recognizes that it is not the only important variable, but does not know what variables are missing. For simplicity, we assume henceforth that the agent conceives of only one situation, "something happens," but knows that her conceptualization is incomplete. In the example, this means that there is only one possible value of x.[4]

At first glance, the standard Savage model seems to provide no way to allow for unforeseen contingencies. Savage takes the set of states of the world as an exogenous element of the model and assumes that the agent has preferences over state-contingent allocations or *acts*. For a model to allow for unforeseen contingencies, it seems necessary to let the state space reflect the agent's subjective understanding of the world, rather than taking it to be exogenous. More precisely, we must identify what the agent believes might happen as a function of her action instead of taking this to be exogenously given.

A natural approach is to think of each possible payoff function as a state, thus constructing a subjective state space. More precisely, if the set of possible actions is A, we could construct a subjective state space where each state is a preference relation over A. This state space seems to be a natural description of how a "fully rational" person should make choices when she is aware that her knowledge of the "true" state space is incomplete. Such an individual does not care about the "real" states per se, caring instead only about how well she does, how she feels as a result of her choice. With this subjective state space in hand, one expects that an individual who is rational in the usual sense would choose in a way that corresponds to forming subjective probabilities over these states and maximizing expected utility.[5] In effect, such a construction would replace unforeseen external possibilities with foreseen payoff possibilities.

However, directly assuming such a state space seems problematic. By analogy, consider the status of subjective probability prior to the work of Savage (1954). Writers like Keynes (1921) and Knight (1921) had advanced strong arguments that behavior under known probabilities (risk) and unknown probabilities (uncertainty) are significantly different. In light of these arguments, the claim that agents under uncertainty would form subjective probabilities and treat them as objective ones would seem quite unconvincing on its own. The importance of Savage (1954) is that he characterized the kinds of preferences (behavior) that

[4] Recently, Ozdenoren (1999) has shown how to extend several of our main results to the case where there is a finite set of objective states or, in the terminology of this paragraph, values of x.

[5] An advantage of defining the state set this way is that the only exogenous element required is the set of feasible actions, A. Kreps' approach also takes such a set as the only exogenous element. A related approach defines the state space from actions A and a set of consequences X as X^A (see Fishburn (1970, Chapter 12.1)).

correspond to having subjective probabilities. This characterization plays two distinct roles. First, one can use it to "justify" the assumption of subjective probability by arguing that the behavior of interest falls into the class corresponding to subjective expected utility. Second, if one believes that risk and uncertainty are different, Savage's characterization may help identify the essential characteristics of preferences for which they do differ. This is precisely the contribution of Ellsberg (1962) who identified the role of the sure-thing principle in ruling out uncertainty-averse behavior, paving the way for new approaches to modeling choice under uncertainty.

Analogously, we believe that simply assuming a subjective state space rules out potentially interesting aspects of unforeseen contingencies by fiat and hence is, at best, a partial answer to the problem. By providing a characterization of preferences that correspond to this model, we hope to clarify the nature of such an assumption. As in the case of Savage, this clarification might be taken as evidence that this approach to modeling unforeseen contingencies is appropriate or as an avenue for identifying less "Bayesian" approaches. At this point in time, it is too early for us to know which interpretation is more appropriate.

There is one difficult question of interpretation: does the agent "foresee" the subjective contingencies that are part of the representation? Normally, we do not worry about such issues. For example, in Savage, we say that the agent behaves as if she had subjective probabilities and do not concern ourselves with the question of whether this, in fact, describes her decisionmaking process. Here, though, the situation is less clear. By assumption, we are representing an agent who *cannot* think of all (external) possibilities with an agent who has a coherent view of all payoff possibilities. If the agent does foresee the payoff possibilities, do we really have unforeseen contingencies? We remain agnostic on this point. The key idea is that we have *allowed for* the possibility of unforeseen contingencies by dropping the assumption of an exogenous state space and characterized the agent's subjective view of what might happen. Whether the agent actually fails to foresee any relevant situations is a different matter. It could be that our representation of the agent is quite literally correct —that is, the agent does in fact foresee the set of future utility possibilities and maximizes as in our representation. In this sense, it is the agent, not the modeller, who replaces unforeseen external possibilities with foreseen utility possibilities when making decisions so that, arguably, he has no truly unforeseen contingencies. On the other hand, as is common in decision theory, one can interpret the model as an "as-if" representation of an agent who cannot imagine the set of situations that might occur.[6] For clarity, we typically refer to *subjective contingencies* and avoid the phrase "unforeseen contingencies."

[6] In the context of contracting, Maskin and Tirole (1999) observe that *if* all relevant utility possibilities are common knowledge, then the fact that physical possibilities may not be known is irrelevant. On the other hand, the "as-if" interpretation of our representation does not seem to permit use of their mechanism, so that their result regarding the irrelevance of unknown physical possibilities is sensitive to this interpretational issue.

1.3. *Background and Detailed Results*

Dropping Savage's assumption of exogenous states requires replacing the objects over which the agent's preferences are defined. Instead of taking an act to be a state-contingent consumption bundle, Kreps viewed an act as determining only the subset of actions, e.g., consumption bundles, from which the agent will subsequently choose. Kreps' showed that if preferences over sets are sufficiently "well behaved," then the agent indeed acts as if she had a subjective state space describing her uncertainty regarding her ex post preferences and is a standard expected-utility maximizer with respect to this uncertainty.

Intuitively, Kreps identifies the agent's view of the possible states of the world from preferences for flexibility in the same way that Savage identified subjective probabilities from preferences over bets. For concreteness, imagine the problem of an agent who must decide now on the menu from which she will have to choose at dinner on a specific night several months away. Let B denote the finite set of (deterministic) options—food items in this example—and consider the agent's preference, $\succeq$, over nonempty subsets of B—called menus—which are denoted as $x \in X = 2^B \setminus \{\varnothing\}$. A choice of a menu is interpreted as a commitment to choose "in the future" from this subset. If the agent knows exactly what her future preference over B will be, say $\succ^*$, we can derive her preference over X from it as follows: If the best (according to $\succeq^*$) element of x is preferred to the best element of x', then $x \succeq x'$. It is easy to see that such a preference over menus will not value flexibility. That is, no preference over menus that is generated in this way can have both $\{b, b'\} \succ \{b\}$ *and* $\{b, b'\} \succ \{b'\}$. In this sense, Kreps argued, it is the desire for flexibility that reveals the agent's uncertainty about her ex post preferences over B. Given our assumption that the agent can only conceive of one possible *exogenous* situation, this means that the agent perceives other *subjective* contingencies.

Turning to the specifics of the representations, under mild axioms, Kreps (1979) derives a representation of preferences over menus, where menu x is evaluated by

$$V(x) = \sum_{s \in S} \max_{b \in x} U(b, s). \tag{1}$$

To understand this representation, imagine that the agent chooses menu x, knowing that at some unmodeled ex post stage, she will learn the state of the world, s, and thus learn her preferences as represented by $U(\cdot, s)$. She then chooses the best object from menu x according to these ex post preferences. Ex ante, these preferences are aggregated by summing the maximum utilities across states. Equivalently, we can think of the states as equally likely and view this as an expectation over s. We refer to this as an *additive representation* since the payoffs are being summed over S. One important point is that S and the $U(\cdot, s)$ functions are part of the representation, not a primitive of the model. In this sense, the model does not assume that the agent foresees all possible future circumstances but yields the conclusion that the agent acts *as if* she had a

coherent view of the possible future utilities. The set S itself is not directly relevant—it is merely an index set. The important aspect of the agent's beliefs is the set of possible ex post preferences, those induced by the collection of utility functions $\{U(\cdot, s) \mid s \in S\}$. We refer to this collection of possible ex post preferences as the *subjective state space*. We use the less specific term *state space* to refer to any convenient index set such as S.

Kreps also considered *ordinal representations*, where $\succeq$ is represented by

$$V(x) = u\left(\left(\max_{b \in x} U(b, s)\right)_{s \in S}\right), \tag{2}$$

where u is some strictly increasing but not necessarily additive function. He showed that the set of preferences with a representation of this form is the same as the set of preferences represented by (1)—thus additivity does *not* impose a restriction on preferences. In either case, the representation is hardly pinned down. To see the point, consider the following example.

Example 1: Suppose $B = \{b_1, b_2, b_3\}$ and the agent's preferences over menus are that she prefers longer menus to shorter. That is, letting $\#x$ denote the number of items in the set x, we have $x \succ x'$ if and only if $\#x > \#x'$. This preference satisfies Kreps' axioms and therefore has an additive representation. In fact, it has several such representations. In particular, consider the subjective state spaces $\bar{S} = \{\bar{s}_1, \bar{s}_2, \bar{s}_3\}$ and $\hat{S} = \{\hat{s}_1, \hat{s}_2, \hat{s}_3\}$ with utility functions $\bar{U}(b, s)$ and $\hat{U}(b, s)$ given by

	$\bar{s}_1$	$\bar{s}_2$	$\bar{s}_3$	$\hat{s}_1$	$\hat{s}_2$	$\hat{s}_3$
b_1	2	1	1	2	1	0
b_2	1	2	1	0	2	1
b_3	1	1	2	1	0	2

With either subjective state space, the function $V(x)$ as defined in (1) gives $V(x) > V(x')$ if and only if $\#x > \#x'$, so each of these V functions represents the preferences. Note that the collection of ex post preferences in $\bar{S}$ is disjoint from the collection in $\hat{S}$: in the latter, there are never any ties in the ex post preferences, while there always are ties in the former. It is also easy to see that the union of these two subjective state spaces also yields an additive representation.

The indeterminacy of the subjective state space is troubling for several reasons.[7] First, it undermines the strength of the conclusion that the agent acts as if she had a coherent subjective state space. By analogy, part of the appeal of Savage (1954) or Anscombe-Aumann (1963) as a justification of subjective

[7] In principle, we may not need to achieve uniqueness. By analogy, in modeling risk, utility functions are only identified up to positive affine transformations, not uniquely, yet the Arrow-Pratt measure of risk aversion is well defined. In fact, Kreps (1979, Theorem 2) characterizes the set of transformations of state spaces that preserve preferences. However, there does not seem to be any simple or useful statement of this set of transformations.

probability is the fact that the subjective probabilities are unique. Second, it clearly causes difficulties in deriving subjective probabilities. It seems impossible to identify the agent's probability distribution on the subjective state space without identifying the latter. In a related vein, applications of the model naturally involve more than one agent. But in a multi-agent extension, we would want to formalize the notions of common knowledge and common priors, which depend on the joint subjective state space. Finally, applications of the model seem to require some measure of the agent's aversion to uncertainty regarding future contingencies, which would presumably be based on the size of S, and, loosely, on the variance of $U(\cdot, s)$ across states. If we cannot identify the subjective state space in a meaningful way, then we have no obvious way to characterize such notions and hence seem unable to use the model effectively. For instance, a natural intuition is that if, after specifying exogenous states as completely as possible, one agent has a larger subjective state space than another, then she is more "averse to subjective contingencies." Yet Example 1 shows that one agent's subjective state space (say, $\bar{S}$) can be a strict subset of another's ($\bar{S} \cup \hat{S}$) while both have the same preferences.

To address this problem, we enrich the choice space by allowing menus of lotteries, instead of considering only menus of deterministic options.[8] To see why this helps, consider Example 1 again. Suppose we take the utilities given in the two subjective state spaces to be von Neumann-Morgenstern utilities. Consider the menus

$$x_1 = \{b_1, (.5)b_2 + (.5)b_3\},$$

where the second item is a lottery giving b_2 with probability $1/2$ and b_3 otherwise, versus

$$x_2 = \{b_1, (.5)b_2 + (.5)b_3, (.5)b_1 + (.5)b_2\}.$$

With the first representation, the payoffs to these are $\bar{V}(x_1) = \bar{V}(x_2) = 5$ so the agent is indifferent between these menus. Note, in particular, that the lottery $(.5)b_2 + (.5)b_1$ is never useful to the agent since she always (at least weakly) prefers b_1 or $(.5)b_2 + (.5)b_3$. On the other hand, with the second representation, there is an ex post preference, namely $\hat{s}_2$, in which the agent strictly prefers

[8] Extending the preferences to sets of lotteries is of interest for other reasons as well. First, it is overly restrictive to assume that menus are chosen in a way that the options are deterministic. For example, while menus of lotteries are artificial in the case where B is a set of food items, presumably, the agent is primarily concerned with the "taste attributes" of the food—the kinds of spices used, the temperature and texture of the food, etc.—rather than the dish itself. It seems quite realistic to suppose that a given dish will correspond to a probability distribution on this space, though admittedly a subjective interpretation of these probabilities is more natural. Second, if one is to apply these preferences, then allowing for uncertainty seems necessary, especially in games, where one would want to allow for mixed strategies and incomplete information. Finally, it is worth noting that the set of lotteries *is* easy to conceptualize and create. In this interpretation, the construction here is analogous to that in Anscombe-Aumann (1963), where preferences are assumed to extend to such objects.

$(.5)b_2 + (.5)b_1$ to either of the other lotteries. As a result, we get $\hat{V}(x_1) = 4.5$, while $\hat{V}(x_2) = 5$, implying that the agent strictly prefers x_2 and hence that these representations no longer reflect the same ex ante preferences. What this illustrates is that if we restrict attention to subjective state spaces where all ex post preferences are expected-utility, we can pin down the subjective state space.

This approach also enables us to identify the aggregator in a sense not possible in Kreps' framework. As mentioned above, additivity of the aggregator is not a restriction in Kreps' model.

Example 2: To see this more concretely, consider the state space $\bar{S}$ and state-dependent preferences $\bar{U}$ in Example 1, but with the nonadditive aggregator $\bar{u}(a_1, a_2, a_3) = a_1 \times a_2 + a_3$. This aggregator is inherently nonadditive since there is no monotone transformations that makes it additive. Nevertheless, the preferences given by $\bar{S}$, $\bar{U}$, and $\bar{u}$ using (2) *do* have an additive representation with $U(b, s)$ given by

	s_1	s_2	s_3
b_1	3	0	4
b_2	0	3	4
b_3	1	1	5

In this sense, Kreps' model is unable to determine whether or not the aggregator is additive.

When we extend Kreps' model to allow the agent to prefer commitment in some ex post circumstances, the lack of identification is still more troubling. A preference for commitment is represented by allowing the agent's ex ante view to differ from her ex post view. More specifically, even though the agent will maximize her ex post utility in each state, ex ante the agent prefers *lower* utility in some states, so the aggregator is decreasing in some ex post utilities. If we modify Kreps' framework to allow such ex ante preferences, the nonuniqueness of the aggregator implies that we cannot identify which ex post states are the ones where the agent wants commitment.

Example 3: For a concrete example, consider two state spaces, $S_1 = \{s_1, s_2, s_3, s_4\}$ and $S_2 = \{s_1, s_2, s_5\}$ where the state-dependent utilities are given by

	s_1	s_2	s_3	s_4	s_5
b_1	0	1	0	1	0
b_2	1	0	0	2	2
b_3	2	0	1	0	1

Consider the preferences generated by the aggregator $u_1(a_1, a_2, a_3, a_4) = 2a_1 + 5a_2 - 3a_3 + a_4$ with state space S_1. It is not hard to show that any additive representation of this preference on S_1 must have coefficients with these signs. However, we also obtain an additive representation of this preference on S_2 with aggregator $u_2(a_1, a_2, a_5) = -a_1 + 6a_2 + 3a_5$ and, again, any additive repre-

sentation on this state space must have coefficients with these signs. Note that state s_1 must enter with a strictly positive coefficient in a representation using S_1, but must enter with a strictly negative coefficient in a representation using S_2. In this sense, we cannot identify the agent's ex ante view of the ex post preference in state s_1.

Note that in both examples, the properties of the aggregator are not pinned down because we can change the state space in a way that requires the aggregator to change. Clearly, then, our identification of the subjective state space can potentially eliminate these problems.

Turning to a more concrete statement of our results, *we define a weak EU representation of an ex ante preference, which takes the form of Kreps' ordinal representation in (2) above with two modifications.* First, each $U(\cdot, s)$ is required to be an expected-utility (affine) function (hence the EU in the name). Second, the conditions on the aggregator u are very weak, weaker than those used by Kreps, even in the case where we assume monotonicity. This class of aggregators includes, e.g., the case where the agent evaluates a menu by the worst possible ex post utility it could yield, a potentially interesting model which is excluded by Kreps' requirements.

We show that *the subjective state space and the aggregator for a weak EU representation are essentially unique.* This result should make applications of this approach easier as it makes it possible to relate the structure of the state space to intuitive properties of preferences. For example, *Theorem 2 shows that if one ex ante preference is more "uncertain" than another, then it must have a larger subjective state space for its weak EU representation. Also, if one preference exhibits a stronger desire for flexibility (commitment) than another, then the aggregator is increasing (decreasing) in more states.*

The significance of these results is highlighted by showing that *weak EU representations exist if and only if preferences satisfy a mild set of conditions.* Other than requiring that $\succ$ be a nontrivial and continuous weak order, the only condition we need is that adding the ability to randomize across menu items does not alter the evaluation of a menu. Hence we conclude that *the subjective state space and aggregator are identified under the expected-utility restriction for a very broad class of ex ante preferences.*

Aside from the pragmatic consideration of the results it yields, there is another reason for restricting attention to ex post preferences that satisfy the expected-utility axioms. *If we are willing to restrict attention to strictly increasing aggregators u, then the expected-utility subjective state space is the smallest possible subjective state space for any representation.* That is, we consider *ordinal EU representations,* like Kreps' ordinal representation (2), but where the ex post preferences are required to be expected-utility preferences. We show that given any ordinal EU representation and any other ordinal representation of the same ex ante preference, the ordinal EU representation's subjective state space has smaller cardinality, strictly so in the finite state-space case. We show that *such ordinal EU representations exist if and only if preferences satisfy monotonicity and a*

weak version of a natural adaptation of the independence axiom in addition to weak order, nontriviality, and continuity.

We also show that *if ex ante preferences satisfy an appropriate (unweakened) version of the standard independence axiom, but not monotonicity, then there is a weak EU representation with an additive aggregator*, similar to (1). Thus we characterize when preferences have a more "standard" representation. Moreover, in contrast to Examples 2 and 3 above, our identification of the aggregator implies that when there exists an additive EU representation every aggregator must be additive (up to a monotone transformation) and the signs of the coefficients are unique. This result does not directly enable us to identify probabilities, but opens the door to doing so as we explain in Section 3.1.

1.4. *Outline and Related Literature*

In the remainder of the introduction, we summarize the relevant literature. In Section 2, we set out the model, definitions, and axioms. In Section 3.1, we demonstrate the uniqueness of the weak EU representation and characterize the preferences that have such a representation. In Section 3.2, we turn to ordinal EU representations and show that the subjective state space identified under the expected-utility restriction is the smallest possible subjective state space. We also identify the preferences that have an ordinal EU representation. In Section 3.3, we characterize the preferences for which an additive EU representation exists, and as a corollary to the uniqueness results, show that additivity and the signs of the coefficient are unique. A sketch of the proofs for the characterization results is in Section 4. Complete proofs, where not contained in the text, are in the Appendix. Some concluding remarks are contained in Section 5.

Our survey of the literature includes only the related decision-theoretic work. For a discussion of epistemic approaches to unforeseen contingencies,[9] see Dekel, Lipman, and Rustichini (1998). Aside from the aforementioned work of Kreps, the only papers we know of that take decision-theoretic approaches to unforeseen contingencies are Ghirardato (1996), Skiadas (1997), and Nehring (1999). All three share our view that with unforeseen contingencies, the agent cannot specify the state space precisely and so can only think in terms of events in the true state space. Ghirardato models this by assuming that the agent views an act as yielding a set of consequences in each event, rather than a single consequence. Thus he gives a generalization of subjective expected utility to acts that are correspondences rather than functions. The representation he derives is a generalization of nonadditive probability models.

Both of the other two papers, like ours, do not assume that there is a given set of consequences, instead deriving what can be interpreted as consequences. Skiadas studies preferences over actions conditional on events and derives a

[9] Other approaches are also possible—see, for example, MacLeod (1996) or Al-Najjar, et al. (1999).

representation where the agent has a subjective utility for each action conditional on each event. Intuitively, this represents the agent's "expectation" of the utility consequences over the unforeseen aspects of a situation. The Kreps approach is similar to Skiadas' approach in that both use the agent's preferences to identify the utility consequences of acts as a function of the state. While Nehring, like us, follows Kreps' approach, he analyzes preferences over acts over menus. That is, like Ghirardato, his acts are functions from events to sets, but unlike Ghirardato and like us, his representation involves an implicit ex post stage at which the agent chooses from the appropriate set. Instead of an expected-utility restriction on the ex post preferences, Nehring restricts attention to ex post preferences in which there are only two (thick) indifference curves—that is, his ex post preferences are represented by a utility function that takes only two values. He gives a uniqueness result given this restriction and, hence, his additivity is meaningful in the same sense as ours. He does not have a minimality result—his subjective state space is typically not minimal.

Another related paper is Gul and Pesendorfer (1999). Motivated by the study of temptation and commitment, they also analyze the sets-of-lotteries model, using the same independence axiom and, for some of their results, the same continuity axiom we use. One of their main results is to characterize what in our terms is an additive EU representation. They replace Kreps' monotonicity axiom with the assumption that a union of two sets is always ranked between the two. In our terminology, this axiom effectively requires the subjective state space to be either a singleton (that is, standard expected utility) or a pair of states where commitment is valued in one of two.[10] We do not require any such axiom and so can have many subjective states, including several in which the agent prefers commitment.

2. PREFERENCES: REPRESENTATIONS AND AXIOMS

Let B be a finite set of K *prizes* and let $\Delta(B)$ denote the set of probability distributions on B. A typical subset of $\Delta(B)$ will be termed a *menu* and denoted x (or $\tilde{x}, x', \bar{x}, y$, etc.), while a typical element of $\Delta(B)$, a *lottery*, will be denoted by β. The agent has a preference relation $\succ$ on the set of nonempty subsets of $\Delta(B)$. We endow this collection with the Hausdorff topology; see Appendix A.1 for precise definitions.

We have in mind an environment where the individual first chooses a menu and at a later stage will choose among the elements of this set, but we do not explicitly model this second choice. To clarify, we refer to the preference $\succ$ over menus as the ex ante preference. As discussed above, the representations we consider include sets of preferences over $\Delta(B)$, interpreted as the possible ex

[10] To see this, let A be a set consisting of a healthy snack and a sweet unhealthy snack and let B contain the healthy snack plus a salty unhealthy snack. If the agent considers two different negative circumstances, namely one in which she would be tempted by a sweet snack and one in which she would be tempted by a salty one, then $A \cup B$ may be strictly worse ex ante than both A and B.

post preferences that will govern the agent's later choice from the menu. We use $\succ^*$ to denote a typical ex post preference over $\Delta(B)$.

2.1. *Representations*

We consider three different notions of a representation of such an ex ante preference. Each of the representations is a triple, consisting of a (nonempty) *state space* S, a state-dependent *utility function* $U: \Delta(B) \times S \to \mathbf{R}$, and an *aggregator* $u: \mathbf{R}^S \to \mathbf{R}$. As explained in the introduction, the idea is that the agent's view of her possible ex post preferences over $\Delta(B)$ are summarized by S and U. The aggregator translates the various possible ex post utility levels from a menu into an ex ante comparison. That is, the preference is represented by

$$V(x) = u\left(\left(\sup_{\beta \in x} U(\beta, s)\right)_{s \in S}\right), \tag{3}$$

the natural generalization of Kreps' ordinal representation.

While we refer to S as the state space, it is just an index set, providing a way to refer to the different ex post preferences over $\Delta(B)$ that are summarized by $U(\cdot, \cdot)$. We refer to the collection of these ex post preferences as the *subjective state space*. Formally, given S, U, and $s \in S$, we define $\succ_s^*$ to be the ex post preference relation over $\Delta(B)$ represented by the utility function $U(\cdot, s)$. That is, $\succ_s^*$ and the subjective state space, $\mathbf{P}(S, U)$, are defined by

$$\beta \succ_s^* \beta' \Leftrightarrow U(\beta, s) > U(\beta', s) \qquad \text{and}$$

$$\mathbf{P}(S, U) = \{\succ_s^* \mid s \in S\}.$$

We focus on *EU representations* in which each $U(\cdot, s)$ is an expected-utility—more precisely, affine—function; that is, for all $s \in S$ and all $\beta \in \Delta(B)$, $U(\beta, s) = \sum_{b \in B} U(b, s)\beta(b)$.

> When S is infinite, certain technical issues arise. These are presented in a smaller font and can be skipped without loss of continuity. In particular, for infinite S, we require a topology on the set of all expected-utility preferences that is discussed in Appendix A.2. In the text, we simply take this topology as given.

Our first objective is to characterize representations of the form (3) for which the subjective state space is unique. Clearly, if we allow the aggregator u to ignore certain states we could never obtain such a uniqueness result since one could add or delete such states freely. Hence we restrict attention to "relevant" subjective states.[11] It is easier to define the relevance of a subjective state—that is, an ex post preference—in terms of the state s in the state space S to which it corresponds, rather than in terms of the ex post preference $\succ_s^*$ directly. In the

[11] Alternatively, we could define a weak EU representation without this requirement and then focus on representations with a minimal state space—that is, a space such that we could not eliminate any states and still have a representation on the remainder. It is not hard to use our arguments to show that the results also hold under this approach.

finite case, a state s is relevant in state space S if there is some comparison of menus for which it is key. That is, there are two menus between which the agent is *not* indifferent, even though they yield the same ex post utility for every subjective state other than s.[12] More precisely, we use the following definition.

DEFINITION 1: Given a representation of the form (3) with $\mathbf{P}(S,U)$ finite, a state $s \in S$ is *relevant* if there exist menus x and x' such that $x \sim x'$ and for any $s' \in S$ with $\succ_s^* \neq \succ_{s'}^*$, $\sup_{\beta \in x} U(\beta, s') = \sup_{\beta \in x'} U(\beta, s')$.

If $\mathbf{P}(S,U)$ is infinite, then state s is relevant if for every neighborhood N of s, there exists x and x' with $x \sim x'$ and such that for all $s' \in S \setminus N$, $\sup_{\beta \in x} U(\beta, s') = \sup_{\beta \in x'} U(\beta, s')$.

The weakest of the three representations we consider is a *weak EU representation*.

DEFINITION 2: A *weak EU representation* of $\succ$ is a nonempty set S, a state-dependent utility function $U: \Delta(B) \times S \to \mathbf{R}$, and an aggregator $u: \mathbf{R}^S \to \mathbf{R}$ such that (i) V as defined in (3) is continuous and represents $\succ$, (ii) each $U(\cdot, s)$ is an expected-utility function, (iii) every $s \in S$ is relevant, and (iv) if $s, s' \in S$, $s \neq s'$, then $\succ_s^* \neq \succ_{s'}^*$.

Part (iv) of the definition is for convenience. It enables us to uniquely refer to a state or "index" s in terms of the corresponding subjective state or ex post preference $\succ_s^*$.

Also, while for notational convenience we define the aggregator as a function on $\mathbf{R}^S$, it is meaningful only on the subspace

$$\mathscr{U}^*(S,U) \equiv \left\{ \left(\sup_{\beta \in x} U(\beta, s) \right)_{s \in S} \mid x \subseteq \Delta(B) \right\}.$$

Henceforth, we omit the S and U arguments from $\mathscr{U}^*$ when it is unlikely to cause confusion.

The next representation we consider, an *ordinal EU representation*, strengthens the weak EU representation by requiring the aggregator u to be strictly increasing. This corresponds to strengthening the requirement that every state be relevant in two ways. First, it gives us the "direction" in which a state must be relevant. More specifically, an increase in ex post utility in any state cannot decrease ex ante utility. Second, it requires that every state *always* be relevant in that a change in the utility in that state *always* "counts."[13] We will also have occasion to consider *ordinal representations* that do not impose the EU requirement.

[12] Note that a state in which the agent is completely indifferent among all lotteries could never be relevant in this sense. Hence this (trivial) ex post preference can never be part of a subjective state space.

[13] For an example of an aggregator that satisfies our requirements for a weak EU representation and is weakly but not strictly increasing, let $u(\cdot)$ be the minimum operator. Clearly, if we increase a vector of ex post utilities without changing the minimum of these utilities, this aggregator does not increase. However, every state (that can achieve the minimum for some menu) is relevant in our sense.

DEFINITION 3: An *ordinal EU representation* is a weak EU representation with an aggregator that is strictly increasing on $\mathscr{U}^*(S,U)$. An *ordinal representation* is a triple (S,U,u) that satisfies all the requirements to be an ordinal EU representation except that U need not be an expected-utility function.[14]

Finally, the last kind of representation we consider is an *additive EU representation* in which u is additive across the vector of maximal ex post utilities. Unlike the ordinal EU case, here we do *not* require monotonicity so we do not restrict the weights on the different states to be positive. Thus the additive EU representation is stronger than the weak EU in a different way than the ordinal EU.

DEFINITION 4: An *additive EU representation* is a weak EU representation[15] such that there exists a finitely additive measure μ on S such that, for all $x \subseteq \Delta(B)$,

$$u\left(\left(\sup_{\beta\in x} U(\beta,s)\right)_{s\in S}\right) = \int_S \sup_{\beta\in x} U(\beta,s)\,\mu(ds).$$

2.2. *Axioms*

The axioms that we consider on the ex ante preference relation are the following. The first three we assume throughout.

AXIOM 1 (Weak Order): $\succ$ *is asymmetric and negatively transitive.*

AXIOM 2 (Continuity): *The strict upper and lower contour sets,* $\{x' \subseteq \Delta(B) \mid x' \succ x\}$ *and* $\{x' \subseteq \Delta(B) \mid x \succ x'\}$, *are open (in the Hausdorff topology).*[16]

AXIOM 3 (Nontriviality): *There is some* x *and* x' *such that* $x \succ x'$.

Our three representations differ primarily in terms of the independence-type condition they require. The condition we use for a weak EU representation says that if we enlarge a menu by allowing the agent to randomize over items on the menu, this expansion has no value or cost to her. This axiom has very little of independence to it, though it clearly is related. Formally, for a menu x, let $\text{conv}(x)$ denote its convex hull.

[14] This definition is not complete as we do not give a topology for the set of all preferences. Hence the notion of a weak representation without the EU restriction is not fully defined. However, the only ordinal representations we will need to consider are either ordinal EU or have finite subjective state spaces where topological considerations play no role.

[15] To be more precise, we must expand the definition of a weak EU representation to make S a measure space and require U to be measurable with respect to this space. Since we make no explicit use of such measurability considerations, we avoid a discussion of the details.

[16] The Hausdorff topology is reviewed in Appendix A.1.

AXIOM 4 (Indifference to Randomization (IR)): *For every menu* $x \subseteq \Delta(B)$, $x \sim \text{conv}(x)$.

The axioms we require for ordinal EU and additive EU representations are much more akin to standard independence axioms. To state these, we first need to define convex combinations. We do this by defining the convex combination of two sets to be the set of pointwise convex combinations. That is, for $\lambda \in [0,1]$, define $\lambda x + (1-\lambda)x' \subseteq \Delta(B)$ to be the set of $\beta'' \in \Delta(B)$ such that $\beta'' = \lambda\beta + (1-\lambda)\beta'$ for some $\beta \in x$ and $\beta' \in x'$ where, as usual, $\lambda\beta + (1-\lambda)\beta'$ is the probability distribution over B giving b probability $\lambda\beta(b) + (1-\lambda)\beta'(b)$.

We first give the stronger axiom since the weaker is more easily understood as a relaxation of it.

AXIOM 5 (Independence): *If* $x \succ x'$, *then for all* $\lambda \in (0,1]$ *and all* $\bar{x}$,

$$\lambda x + (1-\lambda)\bar{x} \succ \lambda x' + (1-\lambda)\bar{x}.$$

This is the usual independence axiom, using the definition above for taking convex combinations.

We now explain the normative appeal of this condition. It is easiest to understand the axiom by breaking it into two parts. To understand the first part, suppose we think of $\lambda x + (1-\lambda)\bar{x}$ not as a convex combination of sets as we have defined it, but instead as a randomization over these menus where the agent gets menu x with probability λ and menu $\bar{x}$ otherwise. (We will justify this interpretation momentarily.) Given this, our axiom is precisely the usual independence axiom and is interpreted in precisely the usual way: the difference between $\lambda x + (1-\lambda)\bar{x}$ and $\lambda x' + (1-\lambda)\bar{x}$ is only in the "λ" event, so the preference between these should be the same as the preference between x and x'.[17]

The key, then, is understanding why a rational agent should view this kind of lottery over sets as equivalent to the convex combination of sets we defined. This interpretation can be thought of as a kind of reduction of compound lotteries together with an assumption that the agent is certain she will satisfy the independence axiom ex post, both normatively appealing notions.

To see this most easily, suppose $x = \{\beta_1, \beta_2\}$ and $\bar{x} = \{\bar{\beta}\}$ and consider how the individual should view the gamble giving x with probability λ and $\bar{x}$ otherwise. The individual knows that whatever menu the gamble gives her ex ante, she will choose her preferred element from that set at the ex post stage. There are two sets of circumstances at the ex post stage: those in which she would choose β_1 over β_2 and those in which she would choose β_2 over β_1. In

[17] Nehring (1999) can be thought of as assuming this form of independence in considering lotteries over sets. (This is not entirely accurate since he considers Savage acts over sets, but he uses the standard Savage axioms to reduce such an act to a lottery over sets.) However, he does not follow our next step of identifying lotteries over sets with our definition of convex combinations of sets, an identification that is at the heart of our independence axiom.

the first case, the randomization over menus effectively gives her $\lambda\beta_1 + (1-\lambda)\bar{\beta}$, while in the second, she effectively receives $\lambda\beta_2 + (1-\lambda)\bar{\beta}$.

Compare this situation to the one where we simply give her the menu $\{\lambda\beta_1 + (1-\lambda)\bar{\beta}, \lambda\beta_2 + (1-\lambda)\bar{\beta}\}$—that is, in place of the lottery, we give her the convex combination of menus. Again, there are clearly two sets of relevant circumstances ex post: those in which she would choose $\lambda\beta_1 + (1-\lambda)\bar{\beta}$ from this menu and those in which she would choose $\lambda\beta_2 + (1-\lambda)\bar{\beta}$. Now suppose that she is sure of one thing: her ex post preference will satisfy the independence axiom. In this case, she knows that the circumstances in which she prefers β_1 to β_2 are exactly those in which she prefers $\lambda\beta_1 + (1-\lambda)\bar{\beta}$ to $\lambda\beta_2 + (1-\lambda)\bar{\beta}$. In other words, both the lottery over menus and the convex combination of menus then effectively give her $\lambda\beta_1 + (1-\lambda)\bar{\beta}$ in those circumstances in which she prefers β_1 to β_2 and $\lambda\beta_2 + (1-\lambda)\bar{\beta}$ in all other circumstances. Hence she should be indifferent between the lottery over menus and the convex combination of the menus.

We emphasize that this is a *normative* argument, relying on the idea that the agent is fully rational except that she does not necessarily know the set of states of the world.[18] None of the argument above requires the agent to understand anything about the circumstances in which she would prefer β_1 to β_2, only to imagine that such circumstances could exist and that her ex post preference in such a situation would satisfy the independence axiom.

Our weaker version of independence requires this implication only for certain menus.

AXIOM 6 (Weak Independence): *If* $x' \subset x$ *and* $x \succ x'$, *then for all* $\lambda \in (0,1]$ *and all* $\bar{x}$,

$$\lambda x + (1-\lambda)\bar{x} \succ \lambda x' + (1-\lambda)\bar{x}.$$

In words, if the addition of $x \setminus x'$ to the menu x' strictly improves it, then adding $\lambda(x \setminus x') + (1-\lambda)\bar{x}$ to $\lambda x' + (1-\lambda)\bar{x}$ must also be a strict improvement.

A natural question to ask is why we do not also require a strict preference implication when $x \prec x'$. The reason is that our only use of this axiom will be in conjunction with monotonicity:

AXIOM 7 (Monotonicity): *If* $x \subseteq x'$, *then* $x' \succeq x$.

In other words, bigger sets are weakly preferred—that is, commitment is never valuable. Obviously, monotonicity implies that both of the preferences to which we refer in Axiom 6 must hold weakly.

The following lemma characterizes the relationships among IR, independence, and weak independence.

[18] In particular, she is indifferent to the timing of the resolution of the objective uncertainty.

LEMMA 1: *If $\succ$ satisfies independence, then it satisfies weak independence. If $\succ$ satisfies weak order, continuity, and weak independence, then it satisfies IR.*

3. IDENTIFYING AND CHARACTERIZING THE REPRESENTATION

3.1. *Uniqueness of the Subjective State Space and Aggregator*

In this subsection, we show that if a weak EU representation with a finite state space exists, then *every* weak EU representation of this ex ante preference has the same subjective state space; that is, the subjective state space is uniquely identified. We also show that such representations exist for a very broad class of preferences; in particular, monotonicity is *not* required.

The uniqueness of the subjective state space in turn implies a form of uniqueness of the aggregator. Because the formal definition is notationally cumbersome, we state the idea here and give the details in Appendix B. Recall that the aggregator is a function from vectors of ex post utilities to an ex ante evaluation. A trivial way to alter the aggregator then is to "relabel" the subjective states—that is, to put the ex post utilities into a vector in a different order and change the aggregator accordingly. Naturally, we will say that two aggregators related in this fashion are essentially the same. A second trivial way to alter the aggregator is to rescale some of the ex post utility functions and to alter the aggregator accordingly. That is, we might replace the state s utility function with twice the original function and then change the aggregator by having it divide this component in half before aggregating it with the other ex post utilities as before. Again, we will say that two aggregators related in this fashion are the same. Finally, the aggregator is only meaningful on those vectors in $\mathbf{R}^S$ that can be generated by some menu. That is, we cannot expect to pin down the aggregator at points outside of $\mathscr{U}^*(S,U)$. When the aggregators for all weak EU representations of a given ex ante preference can be related in this fashion, we say that the aggregator is *essentially unique*.

> If the subjective state space is infinite, it is unique but only up to closure. See Appendix A.2 for details. As a result, uniqueness of the aggregator is further complicated in this case by the fact that we could change the state space in a way that does not change the closure and change the aggregator correspondingly. Again, we view this as an essentially irrelevant change.

THEOREM 1: *A. The ex ante preference $\succ$ has a weak EU representation if and only if it satisfies weak order, continuity, nontriviality, and IR.*

B. If an ex ante preference has a weak EU representation with a finite state space, then all weak EU representations of that preference have the same subjective state spaces. Furthermore, the aggregator is essentially unique.

C. More generally, the closures of $\mathbf{P}(S,U)$ for all weak EU representations of $\succ$ coincide.

PROOF SKETCH: The proof of Theorem 1.A is discussed in Section 4. For the intuition behind Theorem 1.B, suppose we have two weak EU representations

(S, U, u) and (S', U', u'), both with finite subjective state spaces, such that $\mathbf{P} \equiv \mathbf{P}(S, U) \neq \mathbf{P}(S', U') \equiv \mathbf{P}'$. If these subjective state spaces are not the same, then there is some ex post preference, $\succ^*_{s_0}$, contained in, say, $\mathbf{P}'$ and not in $\mathbf{P}$.

For each $s \in S \cup S'$, let $\mathscr{L}_s$ denote a lower contour set for the preference $\succ^*_s$ and let $\overline{U}_s$ denote the associated level of utility. That is, $\mathscr{L}_s$ is the set of points on and "below" the indifference curve associated with utility $\overline{U}_s$ for $\succ^*_s$. Let x denote the intersection of these lower contour sets. Because all of these ex post preferences are expected-utility preferences, we know that these indifference curves are linear. As a consequence, we can always choose the lower contour sets so each coincides with a (nontrivial) part of the boundary of x. (See Figure 1 for example.) Let x' denote the intersection of all these lower contour sets *except* for $\mathscr{L}_{s_0}$. Because this lower contour set formed part of the boundary of x, x' must be strictly larger than x as shown in Figure 1. (In the figure, $x' = x \cup y$.)

For any $s \in S$, consider the value of $\sup_{\beta \in x} U(\beta, s)$. It is easy to see that it cannot exceed $\overline{U}_s$ since every point in x is contained in $\mathscr{L}_s$ and so gives utility less than or equal to $\overline{U}_s$. Also, it cannot be less than $\overline{U}_s$ since we have ensured that the indifference curve for state s associated with this level of utility intersects the boundary of x. Hence for all $s \in S$, this supremum must exactly equal $\overline{U}_s$. Note that exactly the same argument applies to the value of $\sup_{\beta \in x} U'(\beta, s)$ for any $s \in S'$. Also, exactly the same argument applies to x' for any $s \neq s_0$. On the other hand, the argument does not apply to x' for state s_0. Note that x' must contain some points outside $\mathscr{L}_{s_0}$ and, by definition, all such points give utility in state s_0 strictly greater than $\overline{U}_{s_0}$. Hence

$$\sup_{\beta \in x} U(\beta, s) = \sup_{\beta \in x'} U(\beta, s), \quad \forall s \in S,$$

$$\sup_{\beta \in x} U'(\beta, s) = \sup_{\beta \in x'} U'(\beta, s), \quad \forall s \in S', s \neq s_0,$$

and

$$\sup_{\beta \in x} U'(\beta, s_0) < \sup_{\beta \in x'} U'(\beta, s_0).$$

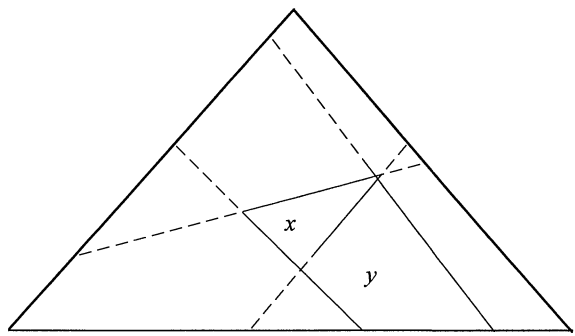

FIGURE 1

Because $s_0 \notin S$, the fact that (S, U, u) represents the ex ante preference $\succ$ implies that $x \sim x'$. Hence in representation (S', U', u'), the aggregator must be ignoring the state s_0 utility difference between x and x'. Roughly, the proof of Theorem 1.B shows that essentially *every* comparison of sets that have utility differing only in state s_0 can be written as a comparison of such an x and x'. Therefore, the aggregator u' must *always* ignore the state s_0 utility difference when it is the only state where a utility difference exists. But this implies that s_0 is not relevant to the representation (S', U', u'), a contradiction.

As noted in the introduction, one reason the lack of identification of the subjective state space in Kreps' framework is problematic is that it makes it difficult to relate the structure of the subjective state space to intuitive properties of the underlying ex ante preferences. For example, a natural intuition is that larger subjective state spaces correspond to a greater concern about subjective contingencies. We also want to give a similar characterization of greater preference for flexibility or commitment using the way subjective states enter the representation. Examples 1 and 3 in the introduction showed that such comparisons cannot be made in the Kreps framework: alternative representations of the *same* preference can have (i) nested subjective state spaces or (ii) identical subjective states with oppositely signed coefficients. Since we pin down the subjective state space (given the EU restriction), we *can* make such comparisons.

To do so, given a weak EU representation (S, U, u) with a finite subjective state space, say that $s \in S$ is *positive* if there are vectors $U^*, \overline{U}^* \in \mathscr{U}^*$ that differ only in coordinate s and have $U_s^* > \overline{U}_s^*$ such that $u(U^*) \geq u(\overline{U}^*)$. Define $s \in S$ to be *negative* if the same is true except that $u(U^*) < u(\overline{U}^*)$. (Note that some states may be both negative and positive.) Let $\mathscr{P}$ denote the set of ex post preferences corresponding to the positive states—that is,

$$\mathscr{P} = \{\succ^* \mid \succ^* = \succ_s^* \text{ for some positive } s\}.$$

Define $\mathscr{N}$ similarly for the negative states. Intuitively, the size of $\mathscr{P}$ measures the agent's desire for flexibility, while the size of $\mathscr{N}$ measures his desire for commitment.

> When the subjective state space is infinite, say that $s \in S$ is *positive* if for every neighborhood N of s, there are menus x and x' with $x \subset x'$, $x' \succ x$, and $\sup_{\beta \in x} U(\beta, s') = \sup_{\beta \in x'} U(\beta, s')$ for all $s' \in S \setminus N$. The definition of a negative state is analogous. Let $\mathscr{P}$ denote the closure of the set of ex post preferences corresponding to positive states and define $\mathscr{N}$ analogously for the negative states.

To relate these attributes of the representation to the preferences, say that agent 2 *desires more flexibility* than agent 1 if $x \cup x' \succ_1 x$ implies $x \cup x' \succ_2 x$ and that agent 2 *desires more restrictions* than agent 1 if $x \cup x' \prec_1 x$ implies $x \cup x' \prec_2 x$. Finally, say that agent 2 *is more uncertain than* agent 1 if $x \cup x' \nsim_1 x$ implies $x \cup x' \nsim_2 x$. Focusing on the intuition of the last condition, $x \cup x' \nsim_1 x$ says that agent 1 conceives of a circumstance in which having the options in $x' \setminus x$ would be relevant to her in some fashion. (For instance, in the monotonic

case, this would mean that she would value the flexibility of having the additional options in $x' \setminus x$.) The implication says that agent 2 must also consider a circumstance in which $x' \setminus x$ is important. In this sense, 2 has more uncertainty than 1. Note that in contrast to the desires for flexibility and commitment, in this last condition there is no requirement that 1 and 2 give the same value to adding x'—that is, we allow $x \cup x' \succ_1 x$ and $x \cup x' \prec_2 x$.

THEOREM 2: *Let (S_i, U_i, u_i) be weak EU representations of preferences $\succ_i$, for $i = 1, 2$.*

1. *If $\succ_2$ desires more flexibility than $\succ_1$, then $\mathscr{P}_1 \subseteq \mathscr{P}_2$.*
2. *If $\succ_2$ desires more restrictions than $\succ_1$, then $\mathscr{N}_1 \subseteq \mathscr{N}_2$.*
3. *If $\succ_2$ is more uncertain than $\succ_1$ and if S_2 is finite, then $\mathbf{P}(S_1, U_1) \subseteq \mathbf{P}(S_2, U_2)$.*
4. *If $\succ_2$ is more uncertain than $\succ_1$ and S_2 is infinite, then the closure of $\mathbf{P}(S_2, U_2)$ must contain the closure of $\mathbf{P}(S_1, U_1)$.*

This result implies that if $\succ_1 = \succ_2$, then $\mathscr{P}_1 = \mathscr{P}_2$ and $\mathscr{N}_1 = \mathscr{N}_2$.[19] Thus we uniquely identify which states are positive and which are negative.

Hence we see that our identification of the subjective state space and aggregator enables us to relate the representation to intuitive, economically meaningful properties of the underlying preferences.

3.2. *Minimality of the EU Subjective State Space*

In this subsection, we give a different reason for focusing on subjective state spaces consisting only of expected-utility preferences: such state spaces are the smallest possible, if we restrict attention to ordinal EU representations.[20] It is easy to see that ordinal EU representations require monotonicity, so, unlike the previous subsection, we do assume monotonicity here.

THEOREM 3: *A. $\succ$ has an ordinal EU representation if and only if it satisfies weak order, continuity, nontriviality, weak independence, and monotonicity.*[21]

[19] This is also an implication of Theorem 1.

[20] This result is *not* true for weak representations, even in the monotonic case. If one assumes that ex ante preferences over menus are generated by having one ex post preference that is not expected-utility but does have convex lower contour sets, the induced preference over menus will have a monotonic weak EU representation. (This follows from Theorem 1.A.) However, this representation will require more than one ex post preference in its subjective state space, while the "correct," non-EU subjective state space is a singleton.

[21] The conditions Kreps uses to prove existence of an ordinal representation are weak order, monotonicity, and $x \sim x \cup x' \Rightarrow x \cup x'' \sim x \cup x' \cup x''$. It is not difficult to show that this last condition is also necessary for an ordinal EU representation in our model, so it must be an implication of our axioms. A direct proof of this fact is not difficult. First, note that $x \cup x' \subseteq x \cup x' \cup x''$, so monotonicity implies that the latter is weakly preferred. By weak independence, then, $\frac{1}{2}x + \frac{1}{2}[x \cup x' \cup x''] \succeq \frac{1}{2}x + \frac{1}{2}[x \cup x'']$, strictly so if $x \cup x' \cup x'' \succ x \cup x'$. Suppose that $x \sim x \cup x'$. Then weak independence implies $\frac{1}{2}x + \frac{1}{2}[x \cup x''] \sim \frac{1}{2}[x \cup x'] + \frac{1}{2}[x \cup x'']$, so $\frac{1}{2}x + \frac{1}{2}[x \cup x' \cup x''] \succeq \frac{1}{2}[x \cup x'] + \frac{1}{2}[x \cup x'']$. However, the left-hand side is a subset of the right-hand side. Hence by monotonicity, we must have indifference. But then this requires $x \cup x' \cup x'' \sim x \cup x'$. We thank Klaus Nehring for showing us a critical step in this argument.

B. If there is an ordinal EU representation with a finite subjective state space, then any ordinal representation of the same ex ante preference that has a different subjective state space must have a strictly larger one.

C. If there is an ordinal EU representation with an infinite subjective state space, then every ordinal representation of the same ex ante preference has an infinite subjective state space. In addition, there must be an ordinal EU representation with a countable subjective state space so every ordinal representation has a subjective state space with weakly larger cardinality.

PROOF SKETCH: We describe the intuition for part B of Theorem 3. Fix any ex post preference, say $\succ^*$, in the subjective state space of an ordinal EU representation, say representation 1, and any (interior) lower contour set, say x, for that preference. Fix any set of lotteries y that is disjoint from x. If the agent's preferences are given by $\succ^*$, then she is strictly better off choosing from $x \cup y$ than from x alone. Since x is the set of lotteries yielding utility less than some amount according to $\succ^*$, everything in y must yield higher utility. Since $\sup_{\beta \in x} U(\beta, s) \leq \sup_{\beta \in x \cup y} U(\beta, s)$ for all s, the fact that the aggregator is strictly increasing implies that $x \cup y \succ x$. Hence if we have another representation of these preferences, say representation 2, this property must be preserved.

How can it be preserved? One way to do so is to ensure that representation 2 contains a preference for which x is a lower contour set. It turns out that if representation 2 also has a finite state space, this is the *only* way to ensure this property. (If 2 has an infinite subjective state space, then our minimality property holds, so we are done in this case.)

In light of this, fix any interior β. For each ex post preference $\succeq_s^*$ in the subjective state space for representation 1, let $\mathscr{L}_s$ denote the lower contour set in which β is maximal, i.e., $\mathscr{L}_s = \{\beta' : \beta \succeq_s^* \beta'\}$. Since these are all expected-utility preferences, each different ex post preference must be associated with a different $\mathscr{L}_s$. But, by the preceding paragraph, each of these lower contour sets must be associated with a different ex post preference in representation 2. Hence representation 2 must have at least as many possible ex post preferences as representation 1. In fact, the proof of Theorem 3.B shows that this comparison must be strict unless the subjective state spaces are the same.

In short, an ordinal EU representation exists if and only if preferences satisfy monotonicity and a weak independence axiom. Such a representation has the minimum possible cardinality over all ordinal representations, EU or otherwise. In this sense, the EU representations have the "simplest" possible subjective state space.

3.3. *Additive Representations*

It is natural to explore when an additive representation exists, as that would be more similar to standard representations of preferences under uncertainty. The next theorem shows that strengthening IR to independence characterizes additivity. Moreover, in contrast to Example 2, the aggregator in any weak EU representation of such preferences must be a monotone transformation of an

affine function, pinning down additivity.[22] Finally, in contrast to Example 3, the uniqueness result of Theorem 1 immediately implies that the signs of the coefficients are uniquely identified.

There is also a novel behavioral implication of independence in this context: in contrast to the general case discussed earlier, if u is affine, then $\mathscr{P} \cap \mathscr{N}$ is empty. That is, any given subjective state either leads to a desire for flexibility *or* for commitment, but not sometimes one and sometimes the other, depending on the sets being compared.

THEOREM 4: *A. The ex ante preference $\succ$ has an additive EU representation if and only if it satisfies weak order, continuity, nontriviality, and independence. If $\succ$ also satisfies monotonicity, then the measure μ is always positive.*

B. When an additive EU representation exists, every weak EU representation has an affine aggregator up to a monotone transformation. That is, if (S, U, u) is a weak EU representation of a preference satisfying these axioms, then there exists a finitely additive measure μ on S such that for any $x \subseteq \Delta(B)$,

$$u\left(\left(\sup_{\beta \in x} U(\beta, s)\right)_{s \in S}\right) = \int_S \sup_{\beta \in x} U(\beta, s)\, \mu(ds)$$

up to a monotone transformation.[23]

Given that we have identified additivity, a natural hope is that we can identify the agent's probabilities over the subjective state space. Unfortunately, this is not straightforward even in the monotonic case. The key to the preference for flexibility is the fact that the agent does not know what his ex post preferences will be. Hence it is critical that ex post preferences vary with the "state of the world." But in the usual Savage or Anscombe-Aumann setting, it is precisely the state independence of preferences that allows identification of probabilities.[24]

To identify probabilities, we must introduce some form of separability in the ex post preferences in such a way that some aspect of state independence is introduced. Since this is a significant deviation from the rest of our work, we

[22] This suggests that a distinction between ordinal and additive representations might be obtained in Kreps' model by means of a restriction on the ex post preferences, analogous to the way we require ex post preferences to be expected-utility. This cannot be done within the class of preferences Kreps studies: we can show by example that there is *no* restriction on ex post preferences in Kreps' framework that (a) allows an additive representation of every ex ante preference he considers and (b) does not allow any intrinsically nonadditive ordinal representation.

[23] This result does *not* just say that given any weak EU representation of preferences that have an additive EU representation, there is an additive EU representation that is a monotone transformation of the weak EU representation. Since any two functions representing the same preferences must be monotone transformations of one another, this would be trivially true. Instead, the result says that this is true without essentially changing the subjective state space.

[24] With state dependent utility, one can always rescale the utility functions to change the probabilities, so the probabilities are meaningless. This does not contradict our result on the essential uniqueness of the aggregator as essential uniqueness allows such affine changes.

only sketch a particularly simple version of the idea here. Returning to the meal-planning example of the introduction, suppose the agent cares about food items (what he has for dinner) and money (how much it costs him). In particular, suppose that while he is uncertain about what he will feel like eating on the night in question, he knows how he will value his money. That is, the agent's utility for money—his degree of risk aversion—is independent of any subjective contingencies. Formally, we rewrite the set B as the product of two finite sets, say Z and M. The elements of Z are interpreted as those choices that are affected by subjective contingencies (food items), while the elements of M (amounts of money) are not. Given a distribution $\beta \in \Delta(Z \times M)$, let β_Z be the marginal on Z and β_M the marginal on M. Focusing on the case where S is finite for simplicity, we could consider a representation of the form

$$\sum_{s \in S} \max_{\beta \in x} [U_Z(\beta_Z, s) + U_M(\beta_M)]\, \mu(s)$$

where U_M and each $U_Z(\cdot, s)$ is an expected-utility function. If such a representation exists, all the results above would apply to identifying the subjective state space and the additivity of the representation. In addition, the fact that s does not appear as an argument in the U_M function enables one to use a straightforward variation of standard results to show that the probability distribution μ is also uniquely identified.

While a representation like this has nice properties, it is not a trivial matter to determine the axioms on preferences that generate it or similar representations. We leave this as a topic for future research.

4. PROOF SKETCH OF THE REPRESENTATION RESULTS

The necessity of the axioms is easily shown in each case.[25] This section is devoted to sketching the sufficiency proofs. Before doing so, we state a useful lemma. Let $\mathrm{cl}(x)$ denote the closure of x (in the Euclidean topology on $\Delta(B)$).

LEMMA 2: *If $\succ$ satisfies weak order and continuity, then for all $x \subseteq \Delta(B)$, $\mathrm{cl}(x) \sim x$.*

Intuitively, the Hausdorff distance between a set and its closure is zero, so continuity in the Hausdorff topology requires the agent to be indifferent between these menus.

[25] The only demonstration that is not completely straightforward is showing that weak independence is necessary for an ordinal EU representation. To see this, assume we have such a representation. Suppose $x \subset x'$ and $x' \succ x$. Fix any $\lambda \in (0, 1]$ and any $\bar{x}$. Because $x \subset x'$, we know that $\sup_{\beta \in x} U(\beta, s) \geq \sup_{\beta \in x'} U(\beta, s)$ for all s. Because $x' \succ x$, we know that we must have at least one s for which this inequality is strict. Since U is an expected-utility function, we have $\sup_{\beta \in \lambda x + (1-\lambda)\bar{x}} U(\beta, s) = \lambda \sup_{\beta \in x} U(\beta, s) + (1 - \lambda)\sup_{\beta \in \bar{x}} U(\beta, s)$ and likewise for x'. Hence $\sup_{\beta \in \lambda x + (1-\lambda)\bar{x}} U(\beta, s) \leq \sup_{\beta \in \lambda x' + (1-\lambda)\bar{x}} U(\beta, s)$ iff $\sup_{\beta \in x} U(\beta, s) \leq \sup_{\beta \in x'} U(\beta, s)$. Hence we see that the inequality before last holds for all s and strictly for some s. Since the representation is an ordinal EU, u is strictly increasing, so this implies $\lambda x + (1 - \lambda)\bar{x} \prec \lambda x' + (1 - \lambda)\bar{x}$.

For the rest of this section, we assume that $\succ$ satisfies the axioms in Theorem 1.A. In light of IR and Lemma 2, we henceforth restrict ourselves to the set of closed, convex, nonempty subsets of $\Delta(B)$, denoted by X.

We begin by establishing the existence of a representation of preferences on X using standard results. We then describe how we transform this into the desired representation.

PROPOSITION 1: *If $\succ$ satisfies weak order and continuity, then there is a $V: X \to \mathbf{R}$ that represents $\succ$, that is,*

$$x \succ x' \quad \textit{iff} \quad V(x) > V(x').$$

V is unique up to monotone transformations and continuous (with respect to the Hausdorff topology).

PROOF: Since $\Delta(B)$ is connected, compact, and metric, the space X is separable (see Theorem 4.5.5, page 51, of Klein and Thompson (1984)) and connected (see Theorem 2.4.6, page 20, of Klein and Thompson). Hence all the conditions of Debreu's theorem (see, for instance, Fishburn (1970, Lemma 5.1, page 62)) are satisfied, giving the desired representation. *Q.E.D.*

We now characterize $V(x)$ for closed and convex $x \subseteq \Delta(B)$, showing that we can write

$$V(x) = u\left(\left(\max_{\beta \in x} U(\beta, s)\right)_{s \in S}\right)$$

for some S and expected utility functions U. It is easy to extend this characterization to all nonempty subsets of $\Delta(B)$ as follows. Fix any $x \subseteq \Delta(B)$ such that $x \notin X$. Since each $U(\cdot, s)$ is an expected utility function,

$$\sup_{\beta \in x} U(\beta, s) = \max_{\beta \in \text{conv}(\text{cl}(x))} U(\beta, s).$$

By IR and Lemma 2, for every $x \subseteq \Delta(B)$, $x \sim \text{conv}(\text{cl}(x))$, so we can define

$$V(x) = V(\text{conv}(\text{cl}(x))) = u\left(\left(\max_{\beta \in \text{conv}(\text{cl}(x))} U(\beta, s)\right)_{s \in S}\right)$$

$$= u\left(\left(\sup_{\beta \in x} U(\beta, s)\right)_{s \in S}\right),$$

completing the extension.

To characterize $V(x)$ for $x \in X$, we first show that we can uniquely identify each element x in X with a function σ_x defined on a subset of $\mathbf{R}^K$, denoted S^K,[26] formally defined below. The function is the *support function*—see Rockafellar (1972, page 28). These functions have the form of $\sigma_x(s) = \max_{\beta \in x} U(\beta, s)$

[26] Recall that K is the number of elements of B.

for $s \in S^K$, where U is linear in β, i.e., $U(\beta, s) = \sum_{b \in B} U(b, s)\beta(b)$. In short, we can uniquely identify each menu in X with a function giving the maximum expected utility from that menu over a certain artificial "state space." This will enable us to rewrite V in terms of these functions and then identify a subset of this artificial state space that is the "real" state space. This construction will generate our weak EU representation. After sketching this construction, we explain how strengthening the axioms enables us to further refine the construction to yield an ordinal or additive EU representation.

For convenience, we write $B = \{b_1, \ldots, b_K\}$. Let $S^K = \{s \in \mathbf{R}^K \mid \sum s_i = 0, \sum |s_i| = 1\}$, and let $C(S^K)$ denote the set of continuous real-valued functions on S^K. For intuition, think of an element of S^K as a possible specification of the von Neumann-Morgenstern utilities for each of the K elements of B. Because such utilities are only identified up to affine transformations, we have two "degrees of freedom" in setting a normalization. For essentially technical reasons, it is convenient to normalize by requiring these utilities to sum to zero and requiring their absolute values to sum to one. We order the functions in $C(S^K)$ pointwise as usual—that is, $\sigma \geq \sigma'$ means $\sigma(s) \geq \sigma'(s)$ for all $s \in S^K$. We now map X into $C(S^K)$, denoting the image of x by σ_x, where for any $s = (s_1, \ldots, s_K) \in S^K$,

$$\sigma_x(s) \equiv \max_{\beta \in x} U(\beta, s) \equiv \max_{\beta \in x} \sum_{i=1}^{K} \beta(b_i) s_i.$$

Let C denote the subset of $C(S^K)$ that σ maps X onto; that is, $C \equiv \{\sigma_x \in C(S^K) \mid x \in X\}$. Finally we define the inverse that maps elements of C into X by

$$x_\sigma \equiv \bigcap_{s \in S^K} \left\{ \beta \in \Delta(B) \;\middle|\; \sum_i \beta(b_i) s_i \leq \sigma(s) \right\}.$$

The following lemma gives two useful properties of the mapping of X to C. First, it is a bijection. Second, it is monotonic in the sense that larger sets (in terms of set inclusion) correspond to larger functions (in the pointwise order).

LEMMA 3: 1. *For all $x \in X$ and $\sigma \in C$, $x_{(\sigma_x)} = x$ and $\sigma_{(x_\sigma)} = \sigma$. Hence σ is a bijection from X to C.*

2. *For all $x, x' \in X$, $x \subseteq x' \Leftrightarrow \sigma_x \leq \sigma_{x'}$.*

PROOF: This is a standard result that follows immediately from the definitions. See, e.g., Clark (1983), Castaing and Valadier (1977), and Rockafellar (1972). *Q.E.D.*

By the first part of this lemma, we can define a function $W: C \to \mathbf{R}$ by $W(\sigma) = V(x_\sigma)$. That is, because each $\sigma \in C$ is associated with a unique $x \in X$, we can define the "utility" of σ to be the utility of the corresponding menu x.

This almost completes the proof. To see why, suppose, for the moment, that we try defining a weak EU representation by setting $S = S^K$, defining

$$U(\beta, s) = U(\beta, (s_1, \ldots, s_K)) = \sum_{i=1}^{K} \beta(b_i) s_i,$$

and letting $u(\cdot) = W(\cdot)$. This last step is not complete since W is only defined on C, that is, on a certain subset of $\mathbf{R}^S$, while u is supposed to be defined on all of $\mathbf{R}^S$. However, we can define u to equal W on C and extend it any way we like to the rest of $\mathbf{R}^S$.

The (S, U, u) so defined satisfies all but one of the properties for a weak EU representation. It is not hard to see that $u[(\sup_{\beta \in x} U(\beta, s))_{s \in S}] = V(x)$ for all $x \in X$. Hence Proposition 1 implies that this function is appropriately continuous and represents the ex ante preference as required. Each $U(\cdot, s)$ is an expected-utility function as it is affine in β.

The only remaining requirement, which will *not* hold in general, is that each $s \in S$ be relevant. Intuitively, S^K includes every ex post preference that we *might* need, but a weak EU representation cannot include ex post preferences that aren't actually needed. Hence we cannot simply set $S = S^K$ but must identify the "relevant" subset of S^K and set S equal to this subset. The proof in the Appendix shows how this can be done.

Let S denote the remaining set of "relevant" points in S^K. Now we can restrict the support functions to this smaller space. The fact that the excluded point was "irrelevant" means that the essence of the construction still works. We lose the bijection property, but if two menus x and x' are associated with the same support function, then we must have $x \sim x'$, so that the utility of the associated support funtion is still well defined. For the rest of this proof sketch, we will continue to use σ to denote a support function but now defined on S instead of all of S^K.

The proof of Theorem 3.A picks up from here, adding the assumptions that $\succ$ satisfies monotonicity and weak independence. We now sketch the proof that these properties imply that the u function constructed above will be strictly increasing on $\mathscr{U}^*$.

This step uses another property of support functions.

LEMMA 4: *For all* $x, x' \in X$, $\sigma_{\lambda x + (1-\lambda)x'} = \lambda \sigma_x + (1 - \lambda)\sigma_{x'}$.

PROOF: This is another standard result. See the same references as for Lemma 3. *Q.E.D.*

We sketch the proof for the case where the weak EU representation identified at the previous step has finitely many states. Recall that one property of weak EU representations is that every state s is relevant. In the finite case, this

requirement is that for every state s, there is an x_1 and x_2 with $x_1 \succ x_2$ but

$$\sup_{\beta \in x_1} U(\beta, s') = \sup_{\beta \in x_2} U(\beta, s')$$

for all $s' \neq s$. So fix any s and such an x_1 and x_2. Without loss of generality, assume $x_2 \subset x_1$. To see why this is without loss of generality, note that if it does not hold, then we can replace x_1 with $x_1 \cup x_2$. By monotonicity, $x_1 \cup x_2 \succeq x_1 \succ x_2$, so we have $x_1 \cup x_2 \succ x_2$. Also, it is easy to see that for all $s' \neq s$,

$$\begin{aligned}\sup_{\beta \in x_1 \cup x_2} U(\beta, s') &= \max\left\{ \sup_{\beta \in x_1} U(\beta, s'), \sup_{\beta \in x_2} U(\beta, s') \right\} \\ &= \sup_{\beta \in x_2} U(\beta, s').\end{aligned}$$

Thus if x_1 does not satisfy these properties, $x_1 \cup x_2$ will, so we can assume that x_1 does.

Let σ_i be the support function of x_i, $i = 1, 2$. Since $x_2 \subset x_1$, we know that $\sigma_1(s) > \sigma_2(s)$. By hypothesis, $\sigma_1(s') = \sigma_2(s')$ for all $s' \neq s$. Let σ_{-s} denote the vector of values of σ_1 and σ_2 for $s' \neq s$. We know that $x_1 \succ x_2$, so $W(\sigma_1(s), \sigma_{-s}) > W(\sigma_2(s), \sigma_{-s})$.

By weak independence, for any $\bar{x}$ and $\lambda \in (0, 1]$, we must have

$$\lambda x_1 + (1 - \lambda)\bar{x} \succ \lambda x_2 + (1 - \lambda)\bar{x},$$

so $W(\sigma_{\lambda x_1 + (1-\lambda)\bar{x}}) > W(\sigma_{\lambda x_2 + (1-\lambda)\bar{x}})$. By Lemma 4, this implies

$$W(\lambda\sigma_1 + (1 - \lambda)\bar{\sigma}) > W(\lambda\sigma_2 + (1 - \lambda)\bar{\sigma})$$

where $\bar{\sigma}$ is the support function of $\bar{x}$. In particular, we could take $\bar{\sigma}$ to be either σ_1 or σ_2, in which case we see that $W(\sigma_1) > W(\lambda\sigma_1 + (1 - \lambda)\sigma_2) > W(\sigma_2)$. It is not hard to strengthen this to show that $W(\lambda\sigma_1 + (1 - \lambda)\sigma_2)$ is strictly increasing in λ. That is, $W(\sigma_s, \sigma_{-s})$ is strictly increasing in σ_s for $\sigma_s \in [\sigma_2(s), \sigma_1(s)]$. What we want to show is that it is strictly increasing in this coordinate everywhere.

Figure 2 describes the situation in two dimensions. We know that W is increasing moving up from the point labeled σ_2 to the point labeled σ_1. It is not hard to show that this requires W to be increasing all along the line through these two points. So consider another possible value for σ_{-s}, say σ'_{-s}. This value corresponds to the vertical line in Figure 2. The argument needed for the boundary points is a little more complex, so let's focus on the case where this line is in the interior. It is easy to see that we can then identify a point like $\bar{\sigma}$ shown in Figure 2 with the property that there is an appropriate convex combination of points on the line through σ_1 and σ_2 that lies on this line, giving the points labeled $\hat{\sigma}_1$ and $\hat{\sigma}_2$ in the figure. By weak independence, $\hat{\sigma}_1$ and $\hat{\sigma}_2$ must be ordered the same way as σ_1 and σ_2. That is, we know that W must be strictly increasing between these two points and hence, just as claimed above, it is strictly increasing as we move up this line as well. Since σ'_{-s} was

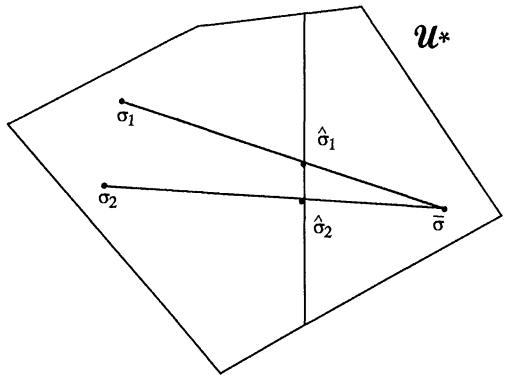

FIGURE 2

(essentially) arbitrary, this implies that W is strictly increasing in the s coordinate everywhere. Since s was arbitrary, this implies that W and hence u is strictly increasing on all of $\mathcal{U}^*$.

Finally, we explain how to develop this construction further in the case when $\succ$ satisfies independence but not necessarily monotonicity to obtain an additive EU representation. First, we amend the very first step of the analysis to strengthen the properties of the V function shown to exist in Proposition 1. The following result, a simple implication of the Herstein and Milnor theorem (see, e.g., Fishburn (1970, Theorem 8.4, page 113), or Kreps (1988, page 54)), is proved in the Appendix (see Section C.5).

PROPOSITION 2: *If $\succ$ satisfies weak order, continuity, and independence, then there is an affine $V: X \to \mathbf{R}$ that represents preferences. That is, the V identified in Proposition 1 can be assumed to satisfy*

$$V(\lambda x + (1-\lambda)x') = \lambda V(x) + (1-\lambda)V(x').$$

V is unique up to affine transformations.

Now that we have added this affinity property of V, it is not hard to see that the W we have constructed will inherit this property. To see why, simply note that by Lemma 4,

$$\sigma_{\lambda x + (1-\lambda)x'} = \lambda\sigma_x + (1-\lambda)\sigma_{x'}.$$

From the way we constructed W, we have

$$V(\lambda x + (1-\lambda)x') = W(\sigma_{\lambda x + (1-\lambda)x'}) = W(\lambda\sigma_x + (1-\lambda)\sigma_{x'}).$$

By the affinity of V, we know that this equals $\lambda V(x) + (1-\lambda)V(x')$. But using the construction of W again, we see that

$$W(\lambda\sigma_x + (1-\lambda)\sigma_{x'}) = \lambda W(\sigma_x) + (1-\lambda)W(\sigma_{x'}).$$

Hence we know that W is continuous and affine. We verify in the Appendix that the structure of C then implies that W can be extended to a continuous linear function on the set of all continuous functions on S. The Riesz representation theorem then implies that W, hence V, can be represented as integrating the value of the function against a measure. Thus there exists μ such that $W(\sigma_x) = \int_{S^K} \sigma_x(s)\mu(ds)$, or $V(x) = \int_{S^K} \sigma_x(s)\mu(ds) = \int_{S^K} \max_{\beta \in x} U(\beta, s)\mu(ds)$, yielding the desired representation.

5. CONCLUSION

To summarize, we have extended Kreps (1979) in several ways. By enriching the structure of the model, we identified an essentially unique subjective state space given a restriction to ex post preferences that are expected-utility preferences. This demonstration is more general than the class of representations Kreps considered, holding for essentially any representation that uses a subjective state space, and is characterized by the property that the agent is indifferent to having the extra option of randomizing over the lotteries in a chosen set. This identification implies that the aggregator is essentially unique as well. In particular, if an additive EU representation exists, then all weak EU representations are additive. As one might expect, additivity is characterized using the independence axiom. None of these results require monotonicity, so we can allow for contingencies in which flexibility is costly. Finally, we showed that in the monotonic case, there is another reason to focus on an EU subjective state space: it is the smallest subjective state space for any ordinal representation. We also showed that ordinal EU representations correspond to preferences satisfying monotonicity and a weak independence axiom.

As illustrated in Theorem 2, pinning down the subjective state space opens up the possibility of giving concrete economic meaning to the properties of the objects in the representation. Our hope is that this paves the way to applications of this model.

In addition to such applications, there are other possible directions for future research. As discussed in Section 3.1, it would be of interest to explore how separability can be used to identify probabilities on the subjective state space. Also, while we have characterized the case where the aggregator is affine, other aggregators might be interesting. For example, perhaps there is an interesting subclass of preferences with a weak EU representation that can be represented using the minimum function as the aggregator.

Finally, as discussed in our introduction, one interpretation of our representation results is that we are determining where an alternative approach must look to find behavior that can be distinguished from this model. If the subjective state space approach misses some interesting aspects of behavior under unforeseen contingencies, it must be true that some axiom (either an explicit one or an implicit assumption built into the structure of the model) precludes this behavior. If there is such an omission, then, just as Ellsberg identified the role of the sure-thing principle in precluding uncertainty-averse behavior, we believe that

one must first find a concrete example of behavior that is a sensible response to unforeseen contingencies but that is precluded by our axioms. An important direction for further research is to see if there is such an Ellsbergian example for this setting and, if so, to explore relaxations of our axioms. We believe that the most interesting possibility is to relax the assumption that the agent knows all the feasible actions. Realistically, part of the problem of unforeseen contingencies is failing to recognize what actions are possible, not just which ones might be useful.

Dept. of Economics, Northwestern University, Evanston, IL 60208, U.S.A., and Tel Aviv University, Tel Aviv, Israel; dekel@northwestern.edu,

Dept. of Economics, University of Wisconsin, 1180 Observatory Dr., Madison, WI 53706, U.S.A.; blipman@ssc.wisc.edu,

and

Dept. of Economics, University of Minnesota, 1035 Heller Hall, 271 19th Ave. S., Minneapolis, MN 55455, U.S.A., and Boston University, Boston, MA, U.S.A.; raldo@bu.edu

Manuscript received November, 1997; final revision received December, 1999.

APPENDIX

A. *Topologies*

A.1. *A Review of the Hausdorff Topology and a Lemma*

Let d denote any distance on $\Delta(B)$. For any pair $x, x' \subseteq \Delta(B)$, we define as usual $d(\alpha, x') \equiv \inf_{\beta \in x'} d(\alpha, \beta)$ and $e(x, x') \equiv \sup_{\alpha \in x} d(\alpha, x')$. The ball in the hemimetric topology is the set defined in (4) below. The topology whose basis is these balls is the Hausdorff hemimetric topology.

$$\mathscr{B}(x, \varepsilon) \equiv \{x' \subseteq \Delta(B) \mid \max\{e(x', x), e(x, x')\} < \varepsilon\}. \tag{4}$$

LEMMA 5: *Let $\{x_n\}$ be an increasing sequence of subsets of $\Delta(B)$, $x_1 \subseteq x_2 \subseteq \cdots$, and let $x^* = \bigcup_{n=1}^{\infty} x_n$. Then $x_n \to \mathrm{cl}(x^*)$ in the Hausdorff topology.*

PROOF OF LEMMA: Fix any $\varepsilon > 0$. Since $\bar{x} \equiv \mathrm{cl}(x^*)$ is compact, there is a finite cover of $\bar{x}$ by open balls of radius $\varepsilon/3$ and center a_m, $m = 1, \ldots, M$. For n^* large enough, x_{n^*} must contains at least one element of each of the M balls, so that $\sup_{a \in \bar{x}} \inf_{b \in x_{n^*}} d(a, b) \leq 2(\varepsilon/3)$. Hence $\lim_{n \to \infty} \sup_{a \in \bar{x}} \inf_{b \in x_n} d(a, b) = 0$, so $x_n \to \bar{x}$. *Q.E.D.*

A.2. *The Topology on* $\mathbf{P}(S, U)$

Let $\mathbf{P}^{EU}$ denote the set of all nontrivial[27] expected-utility preferences. The topology we use on $\mathbf{P}^{EU}$ was also used by Dhillon and Mertens (1996). We define it by specifying which sequences converge. (That this generates a well-defined topology is easily shown using, say, Theorem 2.9 in Kelley (1955).)

[27] That is, $\beta \succ^* \beta'$ for some $\beta, \beta' \in \Delta(B)$.

DEFINITION 5: Given a sequence $\{\succ_n^*\}$ of expected-utility preferences over $\Delta(B)$, we say that $\succ^*$ is a limit of the sequence if it is a nontrivial expected-utility preference such that

$$\beta \succ^* \beta' \quad \text{implies} \quad \exists N \text{ such that } \beta \succ_n^* \beta', \forall n \geq N.$$

Equivalently, we can treat an expected-utility preference as a point in $[0,1]^K$ giving the von Neumann-Morgenstern payoff to each of the K points in B with the normalization that the worst point gets payoff 0 and the best gets payoff 1. The topology above is equivalent to the usual (Euclidean) topology on this space.

Given a subspace $\mathbf{P} \subseteq \mathbf{P}^{EU}$, we define the relative topology on $\mathbf{P}$ in the usual way—that is, $P \subseteq \mathbf{P}$ is open if it is the intersection of $\mathbf{P}$ with an open set in $\mathbf{P}^{EU}$. Given a pair (S,U), we define an induced (relative) topology on S by defining $\hat{S} \subset S$ to be open if $\mathbf{P}(\hat{S},U)$ is open in the relative topology on $\mathbf{P}(S,U)$.

B. *Uniqueness of the Aggregator*

DEFINITION 6: Let $R_i = (S_i, U_i, u_i)$, $i = 1,2$, be weak EU representations of some preferences. If the subjective state spaces of these representations are finite, then R_1 and R_2 are *essentially equivalent* if the following hold.

(i) The subjective state spaces are the same. That is, $\mathbf{P}(S_1,U_1) = \mathbf{P}(S_2,U_2)$.

(ii) There is a bijection $\pi : S_2 \to S_1$ and functions $\gamma : S_2 \to \mathbf{R}_+$ and $\delta : S_2 \to \mathbf{R}$ such that for any $U_1^* \in \mathscr{U}^*(S_1,U_1)$, the vector $g(U_1^*)$ defined by

$$g(U_1^*)(s_2) = \gamma(s_2)U_1^*(\pi(s_2)) + \delta(s_2)$$

is contained in $\mathscr{U}^*(S_2,U_2)$. The function $g : \mathscr{U}^*(S_1,U_1) \to \mathscr{U}^*(S_2,U_2)$ is a bijection.

(iii) Up to a monotonic transformation, $u_1(U_1^*) = u_2(g(U_1^*))$ for all $U^* \in \mathscr{U}^*(S_1,U_1)$.

If the subjective state spaces are infinite, then all of the above holds up to closure. That is, the closures of $\mathbf{P}(S_1,U_1)$ and $\mathbf{P}(S_2,U_2)$ are the same. Also, all references to S_1 and S_2 in (ii) and (iii) are changed to the closures of the set in question.

We can now restate Theorem 1.B: all weak EU representations of a given ex ante preference are essentially equivalent. This follows almost immediately from the uniqueness of the subjective state space, which implies that for every $s_1 \in S_1$, we can find an $s_2 \in S_2$ with $\succ_{s_1}^* = \succ_{s_2}^*$ and vice versa. Given this, it is clear that $U_1(\cdot, s_1)$ must be an affine transformation of $U_2(\cdot, s_2)$ for the corresponding s_2. The g function in the definition simply translates the U^* vectors by rescaling appropriately. The result then follows from the fact that the $V_1(x)$ generated from representation 1 must be a monotone transformation of the $V_2(x)$ generated from representation 2.

C. *Proofs*

For convenience, the order of proofs varies from the order of the results in the text.

C.1. *Proof of Lemma* 2

By definition of $\mathscr{B}(x,\varepsilon)$ (see equation (4) in Appendix A.1), for every $\varepsilon > 0$, $\mathrm{cl}(x) \in \mathscr{B}(x,\varepsilon)$. Recall that $L(x)$ is the strict lower contour set for x. Suppose $x' \succ x$, so that $x \in L(x')$. By continuity, $L(x')$ is open, so by the above, it must be true that $\mathrm{cl}(x) \in L(x')$. That is, $x' \succ \mathrm{cl}(x)$. Similarly, if $x' \prec x$, then $x' \prec \mathrm{cl}(x)$. So suppose $x \succ \mathrm{cl}(x)$. Then $\mathrm{cl}(x) \succ \mathrm{cl}(x)$, contradicting the hypothesis that $\succ$ is a weak order. Similarly, if $\mathrm{cl}(x) \succ x$, we again obtain $\mathrm{cl}(x) \succ \mathrm{cl}(x)$. Hence $x \sim \mathrm{cl}(x)$.[28] *Q.E.D.*

[28] We thank a referee for suggesting this proof, simplifying a proof in a previous draft.

C.2. *Proof of Lemma* 1

The first statement of the lemma is trivially true, so we turn to the second. We show the second by first demonstrating that weak order and weak independence imply that $x \sim \text{conv}(x)$ for every finite x. We then use continuity to complete the argument.

Fix any finite x, let k denote the number of elements of x, and consider the set $\lambda x + (1-\lambda)\text{conv}(x)$ for $\lambda \in (0, 1/k]$. We now show that this set is $\text{conv}(x)$. To see this, note first that $\lambda x + (1-\lambda)\text{conv}(x) \subseteq \text{conv}(x)$ for any λ. For the converse, fix any $\lambda \in (0, 1/k]$ and any $\beta \in \text{conv}(x)$. By definition, there are nonnegative numbers t_i, $i = 1, \ldots, k$, such that $\Sigma_i t_i = 1$ and $\Sigma_i t_i \beta_i = \beta$. Clearly, there must be some j such that $t_j \geq 1/k$. Define $\hat{t}_i$ for $i = 1, \ldots, k$ by

$$\hat{t}_j = \frac{t_j - \lambda}{1 - \lambda}$$

and for $i \neq j$,

$$\hat{t}_i = \frac{t_i}{1 - \lambda}.$$

Obviously, $\hat{t}_i \geq 0$ for all $i \neq j$. Also, $t_j \geq 1/k \geq \lambda$ implies $\hat{t}_j \geq 0$. Finally,

$$\sum_i \hat{t}_i = \frac{1}{1-\lambda}\left[t_j - \lambda + \sum_{i \neq j} t_i\right] = \frac{1}{1-\lambda}[1 - \lambda] = 1.$$

Let $\hat{\beta} = \Sigma_i \hat{t}_i \beta_i$. Clearly, $\hat{\beta} \in \text{conv}(x)$. Hence

$$\lambda \beta_j + (1-\lambda)\hat{\beta} \in \lambda x + (1-\lambda)\text{conv}(x).$$

Clearly, we can write $\lambda \beta_j + (1-\lambda)\hat{\beta} = \Sigma_i t_i' \beta_i$ for some coefficients t_i'. It is easy to see that $t_i' = (1-\lambda)\hat{t}_i = t_i$ for $i \neq j$ and $t_j' = \lambda + (1-\lambda)\hat{t}_j = t_j$. Hence $\lambda \beta_j + (1-\lambda)\hat{\beta} = \beta$. Hence $\lambda x + (1-\lambda)\text{conv}(x) = \text{conv}(x)$.

Of course, $x \subseteq \text{conv}(x)$. Hence weak independence implies that if $x \nsim \text{conv}(x)$, then there is no $\lambda \in [0, 1)$ with $\lambda\,\text{conv}(x) + (1-\lambda)\text{conv}(x) \sim \lambda x + (1-\lambda)\text{conv}(x)$. The left-hand side is $\text{conv}(x)$ and, by the above, there are values of λ for which the right-hand side is $\text{conv}(x)$. Hence $x \sim \text{conv}(x)$ for every finite x.

We now turn to infinite x. By Lemma 2, we can restrict attention to the case where x is closed. Hence x is compact and so has a countable dense subset, say $E = \{e_1, e_2, \ldots\}$. Let $e^n = \{e_1, \ldots, e_n\}$, $n = 1, 2, \ldots$. By the above result, $e^n \sim \text{conv}(e^n)$ for all n. By Lemma 5, $e^n \to \text{cl}(E) = x$ and $\text{conv}(e^n) \to \text{cl}(\text{conv}(E)) = \text{conv}(x)$ in the Hausdorff topology. We now show that this plus $e^n \sim \text{conv}(e^n)$ for all n implies $x \sim \text{conv}(x)$ by continuity. To see this, suppose to the contrary that $x \succ \text{conv}(x)$. Then by continuity, we know that for n sufficiently large, $x \succ \text{conv}(e^n)$ and $e^n \succ \text{conv}(x)$. Fix such an n. By continuity, then, we see that for m sufficiently large, $e^m \succ \text{conv}(e^n)$ and $e^n \succ \text{conv}(e^m)$. But since $\text{conv}(e^n) \sim e^n$, this implies $e^m \succ \text{conv}(e^m)$, a contradiction. The case where $\text{conv}(x) \succ x$ yields a similar contradiction. Hence $x \sim \text{conv}(x)$. *Q.E.D.*

C.3. *Proof of Theorem* 1.*A*

In the text, we gave S^K, $U^K : \Delta(B) \times S^K \to \mathbf{R}$, and $u^K : \mathbf{R}^{S^K} \to \mathbf{R}$ satisfying all the requirements of a weak EU representation except that each $s \in S^K$ be relevant. Recall that s is relevant in state space S if for every neighborhood N (in the relative topology on S) of s, there are menus x and x' with $x \nsim x'$ such that for every $s' \notin N$, $\sigma_x(s') = \sigma_{x'}(s')$. In general, not every $s \in S^K$ will satisfy this requirement.

To construct an appropriate subset, first we define $s \in S^K$ to be *strongly relevant* if for every neighborhood N (in the topology on S^K, not the relative topology) of s, there are menus x and x' in X with $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$ and $x \nsim x'$. Let S^* denote the set of all strongly relevant $s \in S^K$.

Say that $\hat{S} \subseteq S^K$ is *sufficient* if for all x and x' such that $\sigma_x(s) = \sigma_{x'}(s)$ for all $s \in \hat{S}$, we have $x \sim x'$. Clearly, the subjective state space must be sufficient or else it cannot represent the ex ante preference. We will show that S^* is the smallest closed sufficient set. This will be used to show that we can use it for our subjective state space.

First, it is not hard to see that S^* is closed. If $s \in \text{cl}(S^*)$, then every open set containing s intersects S^*. But then each such open set is a neighborhood of some point in S^* so there must be a pair of menus whose support functions differ only on the neighborhood and which are not indifferent. Hence s is also strongly relevant.

We now show that S^* is smaller than every closed sufficient set. Suppose $\hat{S}$ is a closed sufficient set, but $S^* \not\subseteq \hat{S}$. Then there is some $s \in S^*$ with $s \notin \hat{S}$. Because $\hat{S}$ is closed, there is a neighborhood N of s with $N \cap \hat{S} = \varnothing$. By definition of S^*, then, there is an x and x' in X with $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$ and $x \nsim x'$. Hence $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in \hat{S}$ as $\hat{S} \subseteq S^K \setminus N$. Since $x \nsim x'$, this contradicts $\hat{S}$ being sufficient.

In light of the above, we see that if S^* is sufficient, then it is the smallest closed sufficient set. The sufficiency of S^* is an implication of the following lemma. For later use, we prove a result that is more general than is needed here. Let S^*_+ be the set of $s \in S^K$ such that for every neighborhood N of s, there are menus x and x' with $x \subset x'$, $x' \succ x$, and $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$. Intuitively, S^*_+ is the set of positive states—those where flexibility is desirable. Define S^*_- analogously but where $x \succ x'$. Note that $S^*_+ \cup S^*_- \subseteq S^*$. Given any $x, x' \in X$, let

$$D(x, x') = \{s \in S^K \mid \sigma_x(s) \neq \sigma_{x'}(s)\}.$$

LEMMA 6: *If $x \nsim x'$, then $D(x, x') \cap S^* \neq \varnothing$. If $x \subset x'$ and $x' \succ x$, then $D(x, x') \cap S^*_+ \neq \varnothing$. If $x \subset x'$ and $x \succ x'$, then $D(x, x') \cap S^*_- \neq \varnothing$.*

PROOF OF LEMMA: First, note that if we have menus x and x' with $x \nsim x'$ and $D(x, x') \cap S^* = \varnothing$, then without loss of generality, we can assume these sets are nested. To see this, suppose that neither set is contained in the other. Because $x \nsim x'$, at least one of these sets is not indifferent to $\text{conv}(x \cup x')$. Hence we may as well assume that $x \nsim x \cup x'$. Note that if $\sigma_x(s) = \sigma_{x'}(s)$, then $\sigma_x(s) = \max\{\sigma_x(s), \sigma_{x'}(s)\} = \sigma_{\text{conv}(x \cup x')}(s)$. Hence $D(x, x \cup x') \subseteq D(x, x')$. So since $D(x, x') \cap S^* = \varnothing$, the same is true of $D(x, x \cup x')$. Therefore, we may as well assume $x \subset x'$.

Hence it is sufficient to prove the results claimed for S^*_+ and S^*_- as this will imply the claim about S^*. We give the proof for S^*_+; the argument for S^*_- is analogous.

So suppose we have x and x' with $x \subset x'$, $x' \succ x$, and $D(x, x') \cap S^*_+ = \varnothing$. Without loss of generality, we can assume that both x and x' have nonempty interiors. To see this, suppose that one or both have empty interiors. For $\lambda \in (0, 1)$, define $\bar{x}(\lambda) = \lambda x + (1 - \lambda)\Delta(B)$ and define $\bar{x}'(\lambda)$ analogously using x'. It is easy to see that $\bar{x}(\lambda)$ and $\bar{x}'(\lambda)$ have nonempty interiors for all $\lambda > 0$. Also, by Lemma 4, it is easy to see that $D(x, x') = D(\bar{x}(\lambda), \bar{x}'(\lambda))$ for all $\lambda > 0$. By continuity of $\succ$, there exists $\lambda > 0$ such that $\bar{x}(\lambda) \prec \bar{x}'(\lambda)$. Hence if one or both of x and x' have empty interiors, we can replace them with $\bar{x}(\lambda)$ and $\bar{x}'(\lambda)$ for λ sufficiently small. So we may as well assume they have nonempty interiors.

Consider the family of sets $\hat{x} \in X$ such that the following three conditions hold. First, $x \subseteq \hat{x} \subseteq x'$. Second, $\hat{x} \preceq x$. Finally, $D(\hat{x}, x') \subseteq D(x, x')$. It is easy to see that this collection of sets is nonempty as x itself satisfies these conditions. Suppose we have an increasing chain of sets in this family $x_1 \subset x_2 \subset \cdots$. Let x_* denote the closure of the limit of this sequence. We claim that x_* also satisfies the three properties and is in X. To see this, note that the first property is trivially satisfied. By continuity of preferences, the second holds as well: if we have a sequence $\{x_n\}$ with $x_n \preceq x$ for all n, then continuity of $\succ$, the fact that $x_n \to \bigcup_{k=1}^{\infty} x_k$ in the Hausdorff topology (see Lemma 5), and Lemma 2 imply that $x_* \preceq x$. To see that the third property is satisfied, first note that if $\hat{x} \subset \hat{x}'$, then $D(\hat{x}, \hat{x}')$ is the set of s such that

$$\beta \cdot s < \sigma_{\hat{x}'}(s), \quad \forall \beta \in \hat{x}.$$

So consider any $s \in D(x_*, x')$. Since $x_* \subseteq x'$, we have $\beta \cdot s < \sigma_{x'}(s)$ for all $\beta \in x_*$. Since $x \subseteq x_*$, the same is true of all $\beta \in x$, so $s \in D(x, x')$. Hence $D(x_*, x') \subseteq D(x, x')$. Since x_* is closed and convex, we have $x_* \in X$.

Hence any increasing chain of sets in this collection has an upper bound that is also in this collection. Therefore, by Zorn's lemma, this collection of sets has at least one maximal element. Let x^* denote any such maximal element. That is, x^* satisfies the above three properties and there is no $\bar{x}$ strictly containing x^* that does so. We now derive a contradiction by showing that such an $\bar{x}$ must exist.

Recall that x and hence x^* have a nonempty interior and that x^* is closed and convex (since it is an element of X). Hence by Theorem V.9.8 of Dunford and Schwartz (1958), there is a dense subset of the boundary of x^* such that x^* has a unique tangent at each point in this set. Let τ denote such a dense set. Because $x^* \precsim x \prec x'$, we have $x^* \neq x'$. Hence $x^* \subset x'$, implying that there exists a $\beta^* \in \tau$ such that β^* is in the interior of x'.[29] Fix any such β^*. Because $\beta^* \in \tau$, we know that there is a unique $s \in S^K$ such that $\sigma_{x^*}(s) = \beta^* \cdot s$. Let s^* denote this s. Because β^* is in the interior of x', we know that $\sigma_{x^*}(s^*) = \beta^* \cdot s^* < \sigma_{x'}(s^*)$. Hence $s^* \in D(x^*, x')$. Another implication of β^* being in the interior of x' is that there is a sequence $\{\beta_n\}$ such that $\beta_n \in x' \setminus x^*$, $\beta_n \cdot s^* > \beta^* \cdot s^*$, and $\lim_{n \to \infty} \beta_n = \beta^*$. Fix such a sequence and let $x_n = \text{conv}(x^* \cup \{\beta_n\})$. It is not hard to see that $x_n \in X$ and

$$D(x^*, x_n) = \{s \in S^K \mid \beta_n \text{ is unique } \beta \in x_n \text{ such that } \beta \cdot s = \sigma_{x_n}(s)\}.$$

To see why this holds, note that $x^* \subset x_n$ implies $D(x^*, x_n)$ is the set of s such that $\beta \cdot s < \sigma_{x_n}(s)$ for all $\beta \in x^*$. Clearly, $\beta_n \notin x^*$ implies that if β_n is the unique $\beta \in x_n$ such that $\beta \cdot s = \sigma_{x_n}(s)$, then $\sigma_{x^*}(s) < \sigma_{x_n}(s)$, so $s \in D(x^*, x_n)$. If $\beta_n \cdot s < \sigma_{x_n}(s)$, then $\beta \cdot s < \sigma_{x_n}(s)$ for all $\beta \in x_n \setminus x^*$, so $\sigma_{x^*}(s) = \sigma_{x_n}(s)$. Finally, suppose $\beta_n \cdot s = \sigma_{x_n}(s)$ but there is a $\beta' \neq \beta_n$, $\beta' \in x_n$, such that $\beta' \cdot s = \beta_n \cdot s$. Obviously, if $\beta' \in x^*$, then $\sigma_{x^*}(s) = \sigma_{x_n}(s)$. If $\beta' \notin x^*$, then there must be some $\overline{\beta} \in x^*$ and $\lambda \in (0, 1)$ such that $\beta' = \lambda \beta_n + (1 - \lambda)\overline{\beta}$. But then $\beta' \cdot s = \lambda \beta_n \cdot s + (1 - \lambda)\overline{\beta} \cdot s$, so $\beta' \cdot s = \beta_n \cdot s$ implies that there is a $\beta \in x^*$ with $\overline{\beta} \cdot s = \beta_n \cdot s$. Hence, again, $\sigma_{x^*}(s) = \sigma_{x_n}(s)$.

By $\beta_n \cdot s^* > \beta^* \cdot s^* = \sigma_{x^*}(s^*)$, we see that $s^* \in D(x^*, x_n)$ for all n. On the other hand, consider any $s \neq s^*$. For such an s, $\beta^* \cdot s < \sigma_{x^*}(s)$ since s^* is the only s with $\beta^* \cdot s = \sigma_{x^*}(s)$. Hence for n sufficiently large, β_n will be close enough to β^* that $\beta_n \cdot s < \sigma_{x^*}(s)$, implying that $s \notin D(x^*, x_n)$. Hence

$$\bigcap_{n=1}^{\infty} D(x^*, x_n) = \{s^*\}.$$

So for n large enough, we can make $D(x^*, x_n)$ an arbitrarily small set containing s^*.

Recall that $D(x, x') \cap S_+^* = \varnothing$. Hence since $s^* \in D(x^*, x') \subseteq D(x, x')$, $s^* \notin S_+^*$. Therefore, there is a neighborhood N of s^* with the property that for every $\hat{x}$ and $\hat{x}'$ with $\hat{x} \subset \hat{x}'$ and $D(\hat{x}, \hat{x}') \subseteq N$, $\hat{x} \succsim \hat{x}'$. Note that if this is true for N, it is also true for any open subset of N that contains s^*. Hence without loss of generality, we can assume $N \subseteq D(x^*, x')$. Also, by the above, for n sufficiently large, $D(x^*, x_n)$ must be an open subset of N containing s^*. Fix such an n.

By construction, $x \subseteq x^* \subset x_n \subset x'$. Because $D(x^*, x_n) \subseteq N$, $x_n \precsim x^* \precsim x$ so $x_n \precsim x$. By construction, $D(x^*, x_n) \subseteq N \subseteq D(x^*, x') \subseteq D(x, x')$. Hence for all $s \notin D(x, x')$, $\sigma_{x'}(s) = \sigma_{x^*}(s) = \sigma_{x_n}(s)$. Hence $D(x_n, x') \subseteq D(x, x')$. Therefore, we see that x_n satisfies the same properties as x^* and is strictly larger than x^*. Since x^* was the maximal set satisfying these conditions, we have a contradiction. *Q.E.D.*

[29] To see this, note that the denseness of τ implies that we only need to verify that there is a point in the boundary of x^* and the interior of x'. Because a closed convex set equals the convex hull of its boundary, the boundary of x^* cannot equal the boundary of x'. Hence there is some β' in the boundary of x' that is not in x^*. Fix any β'' in the interior of x^*. Because x^* is closed, there is a largest λ such that $\lambda \beta' + (1 - \lambda)\beta'' \in x^*$. The point so defined must be in the boundary of x^* and the interior of x'.

This implies S^* is sufficient. If not, there is some x and x' with $x \nsim x'$ and $S^* \cap D(x, x') = \varnothing$, contradicting Lemma 6. Given the facts shown above, then, S^* is the smallest closed sufficient set.

To complete the construction of the weak EU representation, then, let $S = S^*$. By nontriviality and Lemma 6, we know that $S \neq \varnothing$. For each x, let $\overline{\sigma}_x$ denote the restriction of σ_x to S. Let $\overline{C}$ denote the set of these restricted support functions. By the definition of sufficiency, if $\overline{\sigma}_x = \overline{\sigma}_{x'}$, then we must have $x \sim x'$. Hence there is no ambiguity in defining $\overline{W}: \overline{C} \to \mathbf{R}$ by $\overline{W}(\overline{\sigma}_x) = W(\sigma_x)$. Define $U(\beta, s)$ by restricting U^K to $\Delta(B) \times S$. Finally, let $u(\cdot) = \overline{W}(\cdot)$ on $\overline{C}$ and extend to the rest of $\mathbf{R}^S$ in any fashion. This is a weak EU representation as long as every $s \in S$ is relevant.

To show that every $s \in S$ is relevant, suppose $s \in S$ is not relevant. That is, there is a set N that is open in the relative topology and contains s such that for every x and x' with $\overline{\sigma}_x(s') = \overline{\sigma}_{x'}(s')$ for all $s' \in S \setminus N$, we have $x \sim x'$. But then $S \setminus N$ is sufficient. Recall that S is closed. Since N is open in the relative topology, it equals $S \cap N'$ for some open N'. Hence $S \setminus N = S \setminus N'$ is closed. Hence $S \setminus N$ is a closed sufficient set strictly contained in S, contradicting the minimality of S. *Q.E.D.*

C.4. *Proof of Theorem 3.A*

Given a representation (S, U, u) and a set x, let $U^*(x) \in \mathbf{R}^S$ be defined by $U^*(x)(s) = \sup_{\beta \in x} U(\beta, s)$.

LEMMA 7: *Suppose $\succ$ has a weak EU representation (S, U, u). For any $U^* \in \mathscr{U}^*$, define*

$$x(U^*) = \bigcap_{s \in S} \{\beta \in \Delta(B) \mid U(\beta, s) \leq U_s^*\}.$$

Then $U^*(x(U^*)) = U^*$.

PROOF: Since $U^* \in \mathscr{U}^*$, there must be some $x \in X$ with $(\max_{\beta \in x} U(\beta, s))_{s \in S} = U^*$. Fix any $\beta \in x$. Clearly, $U(\beta, s) \leq \max_{\beta' \in x} U(\beta', s) = U_s^*$ for all s. Hence $\beta \in x(U^*)$. So $x \subseteq x(U^*)$. Therefore, for every s,

$$\max_{\beta \in x} U(\beta, s) \leq \max_{\beta \in x(U^*)} U(\beta, s).$$

But we cannot have $\max_{\beta \in x(U^*)} U(\beta, s) > \max_{\beta \in x} U(\beta, s) = U_s^*$ since no β with $U(\beta, s) > U_s^*$ can be contained in $x(U^*)$ by definition. Hence $\max_{\beta \in x(U^*)} U(\beta, s) = U_s^*$ for all s. *Q.E.D.*

We now add the assumption that $\succ$ satisfies weak independence and monotonicity and show that this implies that the u constructed in the proof of Theorem 1.A is strictly increasing on the appropriate subset of $\mathbf{R}^S$, so that this representation is an ordinal EU representation. Because we are entirely concerned with properties of u on $\mathscr{U}^*$, it is convenient to begin by translating two key facts into statements about u.

First, fix any $s \in S$. Since s must be relevant, we know that for every neighborhood N of s, there are menus x and x' such that $x \succ x'$ and

$$\sup_{\beta \in x} U(\beta, s') = \sup_{\beta \in x'} U(\beta, s') \quad \forall s' \in S \setminus N.$$

Without loss of generality, we can assume $x' \subseteq x$. If this is not true, we can replace x with $x \cup x'$. To see this, note that by monotonicity, $x \cup x' \succeq x \succ x'$, so $x \cup x' \succ x'$. Also,

$$\sup_{\beta \in x \cup x'} U(\beta, s') = \max\left\{ \sup_{\beta \in x} U(\beta, s'), \sup_{\beta \in x'} U(\beta, s') \right\} = \sup_{\beta \in x'} U(\beta, s') \quad \forall s' \in S \setminus N.$$

Hence $x \cup x'$ satisfies the same properties as x, so we can use it instead if $x' \not\subset x$. So we have $x' \subset x$, implying $U^*(x) > U^*(x')$, and we have $u(U^*(x)) > u(U^*(x'))$. Because every open subset of S is a neighborhood of each of its points, this implies that for every open N, there exists[30] $U, U' \in \mathscr{U}^*$ with $U > U'$, $U(s) = U'(s)$ for every $s \notin N$, and $u(U) > u(U')$.

[30] In what follows, we abuse notation slightly by using U to denote a vector of utilities in $\mathscr{U}^*$ instead of the utility function.

Next, in the notation of Lemma 7, if $U > U'$, then $x(U') \subset x(U)$. Hence Lemma 7 and weak independence imply that if $U > U'$ and $u(U) > u(U')$, then for any menu $\hat{x}$,

$$\lambda x(U) + (1-\lambda)\hat{x} \succ \lambda x(U') + (1-\lambda)\hat{x}$$

or, by Lemma 4,

$$u(\lambda U + (1-\lambda)U^*(\hat{x})) > u(\lambda U + (1-\lambda)U^*(\hat{x})).$$

Since $\hat{x}$ is arbitrary, if $U > U'$ and $u(U) > u(U')$, then for all $\lambda \in (0,1]$ and all $\hat{U} \in \mathscr{U}^*$,

$$u(\lambda U + (1-\lambda)\hat{U}) > u(\lambda U' + (1-\lambda)\hat{U}). \tag{5}$$

To complete the proof, we show that for any $U_1, U_2 \in \mathscr{U}^*$ with $U_1 > U_2$, there is a $\tilde{U} \in \mathscr{U}^*$ such that either $U_1 > \tilde{U}$ and $u(\tilde{U}) > u(U_2)$ or $u(U_1) > u(\tilde{U})$ and $\tilde{U} > U_2$. By monotonicity, this will establish that whenever $U_1 > U_2$, we must have $u(U_1) > u(U_2)$.

So fix any $U_1, U_2 \in \mathscr{U}^*$ with $U_1 > U_2$. First, suppose U_2 is in the interior of $\mathscr{U}^*$. Consider the set

$$\{s \in S \mid U_1(s) - U_2(s) > \varepsilon\}$$

for $\varepsilon > 0$. Clearly, there must be an $\varepsilon > 0$ sufficiently small that this set is nonempty. Fix such an ε and let N denote the set above. It is easy to see that N must be open. Hence there are $\hat{U}_1, \hat{U}_2 \in \mathscr{U}^*$ differing only on N such that $\hat{U}_1 > \hat{U}_2$ and $u(\hat{U}_1) > u(\hat{U}_2)$.

Because U_2 is in the interior of $\mathscr{U}^*$ and because $\mathscr{U}^*$ is convex, there must be a $\lambda \in (0,1)$ and $\overline{U} \neq U_2$ such that $\lambda\hat{U}_2 + (1-\lambda)\overline{U} = U_2$. Furthermore, we can choose $\lambda > 0$ arbitrarily small and still find a $\overline{U} \neq U_2$ satisfying $\lambda\hat{U}_2 + (1-\lambda)\overline{U} = U_2$.

In light of this, choose any λ strictly between 0 and the smaller of 1 and

$$\frac{\varepsilon}{\sup_{s \in N}[\hat{U}_1(s) - \hat{U}_2(s)]}.$$

Fix the associated $\overline{U}$ and define $\tilde{U} = \lambda\hat{U}_1 + (1-\lambda)\overline{U}$. By (5), $u(\tilde{U}) > u(U_2)$. Hence we are done with this case if $U_1 > \tilde{U}$. For any s, $U_1(s) \geq \tilde{U}(s)$ iff

$$U_1(s) - U_2(s) \geq \lambda\hat{U}_1(s) + (1-\lambda)\overline{U}(s) - U_2(s) \quad \text{or}$$

$$U_1(s) - U_2(s) \geq \lambda[\hat{U}_1(s) - \hat{U}_2(s)].$$

For $s \notin N$, the right-hand side is zero, so $U_1 > U_2$ implies this. For $s \in N$, we have

$$U_1(s) - U_2(s) > \varepsilon \geq \varepsilon \frac{\hat{U}_1(s) - \hat{U}_2(s)}{\sup_{s' \in N}[\hat{U}_1(s') - \hat{U}_2(s')]} > \lambda[\hat{U}_1(s) - \hat{U}_2(s)],$$

so $U_1 > \tilde{U}$.

The case where U_1 is in the interior of $\mathscr{U}^*$ is completely analogous and so is omitted. Hence we are done if we can show the result for the case where both U_1 and U_2 are in the boundary of $\mathscr{U}^*$. To handle this case, note that monotonicity already implies $u(U_1) \geq u(U_2)$, so if the result does not hold, we must have $u(U_1) = u(U_2)$. So suppose this is true. Fix any interior $\overline{U}$ and $\lambda \in (0,1)$ and define $U_i' = \lambda U_i + (1-\lambda)\overline{U}$, $i = 1, 2$. Clearly, $U_1 > U_2$ implies $U_1' > U_2'$ and both points are in the interior of $\mathscr{U}^*$. By weak independence,[31] $u(U_1) = u(U_2)$ implies $u(U_1') = u(U_2')$. But the argument above applied to U_1' and U_2' shows that we must have $u(U_1') > u(U_2')$, a contradiction. *Q.E.D.*

[31] The usual proof that the "strict form" of independence ($p \succ q$ implies $\lambda p + (1-\lambda)r \succ \lambda q + (1-\lambda)r$) implies the same for indifference (that is, $p \sim q$ implies $\lambda p + (1-\lambda)r \sim \lambda q + (1-\lambda)r$) is easily used to show this. Because U_1 corresponds to a superset of U_2, $\lambda U_1 + (1-\lambda)\overline{U}$ will correspond to a superset of $\lambda U_2 + (1-\lambda)\overline{U}$ for any λ and $\overline{U}$. Hence one can use weak independence exactly the way independence is used in the usual proof.

C.5. *Proof of Theorem* 4.*A*

We now assume that $\succ$ satisfies independence and drop the assumption that it satisfies monotonicity. Instead of continuing with the construction above, it is more convenient to demonstrate additivity on all of S^K before removing the "irrelevant" states.

We first state another useful property of support functions, the proof of which can be found in the same references as given for the proof of Lemma 3.

LEMMA 8: *For all* $x, x' \in X$, $d_{\text{Hausdorff}}(x, x') = d_{\text{supnorm}}(\sigma_x, \sigma_{x'})$.

Next we prove Proposition 2. We do so by verifying that the mixture space axioms (see Kreps (1988, page 52)) hold for X. The only mixture space condition that is not trivial to verify is the Herstein-Milnor continuity condition. We now show that our continuity condition implies that if $x, x', x'' \in X$ and $x \succ x' \succ x''$, then there is a $\lambda_1 \in (0,1)$ and $\lambda_2 \in (0,1)$ such that $\lambda_1 x + (1-\lambda_1)x'' \succ x' \succ \lambda_2 x + (1-\lambda_2)x''$. To see this, let $\lambda x + (1-\lambda)x'' \equiv x(\lambda)$. Then for any $\lambda, \mu \in [0,1]$,

$$d_{\text{Hausdorff}}(x(\lambda), x(\mu)) = d_{\text{supnorm}}(\sigma_{x(\lambda)} - \sigma_{x(\mu)})$$
$$= |\lambda - \mu| d_{\text{supnorm}}(\sigma_x - \sigma_{x''})$$

where the first equality is Lemma 8, and the second follows from Lemma 4. Hence the function from $[0,1]$ to X with the Hausdorff topology defined by $x(\lambda)$ is continuous. Now the result follows from continuity of $\succ$.

REMARK 1: The restriction to X is needed for the mixture space axioms because $\lambda[\lambda' x + (1-\lambda')x'] + (1-\lambda)x'$ might not equal $\lambda\lambda' x + (1-\lambda\lambda')x'$ if x and x' are not convex.

So by the Herstein-Milnor theorem, there is an affine V that represents the preferences and is unique up to an affine transformation. By this uniqueness and Lemma 1, V is continuous in the Hausdorff topology, completing the proof of Proposition 2.

LEMMA 9: 1. *C is convex.*

2. *The zero function is in C, in particular $\sigma_{\{(1/K,\ldots,1/K)\}}(s) = 0$ for all s.*

3. *There exists $c > 0$ such that the constant function equal to c is in C. That is, $\sigma^c \in C$, where $\sigma^c(s) = c$ for all s.*

4. *The supremum of any two elements in C is in C: $\sigma \in C$ and $\sigma' \in C \Rightarrow \sigma \vee \sigma' \in C$, where $(\sigma \vee \sigma')(s) = \max\{\sigma(s), \sigma'(s)\}$.*

PROOF: 1. Given σ_x and $\sigma_{x'}$ in C, using Lemma 4 and the convexity of X, any convex combination of σ_x and $\sigma_{x'}$ is in C.

2. For any $s \in S^K$, we have $\Sigma_i s_i = 0$ so by definition $\sigma_{\{(1/K,\ldots,1/K)\}}(s) = \Sigma_i (1/K)s_i = 0$.

3. First note that $\sigma_{\Delta(B)}(s) = \max_i\{s_i\} \geq 1/(2K)$. The equality follows from the definition of the support functions. The inequality follows from the definition of S^K. (If $\max_i\{s_i\} < 1/(2K)$, then $\Sigma_{\{i \mid s_i > 0\}} s_i < 1/2$. Then, since $\Sigma_i s_i = 0$, also $\Sigma_{\{i \mid s_i < 0\}} |s_i| < 1/2$. But then, $\Sigma_i |s_i| < 1$, which contradicts the definition of S^K.) Now consider $x \equiv \{\beta \mid \Sigma_i \beta(b_i)s_i \leq c$ for all $s \in S^K\}$, where $0 < c < 1/(2K)$. Clearly x is a closed, convex, and nonempty subset of $\mathbf{R}^K$. It is easy to see that we could have defined our mapping from X into $C(S^K)$ to have as its domain all convex, closed nonempty subsets of $\mathbf{R}^K$ without affecting any of our lemmas on support functions. With this definition, clearly, σ_x is the constant function c. It remains to show that $x \in X$. By part (2) of Lemma 3, since $\sigma_x \leq \sigma_{\Delta(B)}$ we know that $x \subseteq \Delta(B)$, so $x \in X$.

4. Given σ_x and $\sigma_{x'}$ in C, it is easy to see that $\sigma_{\text{conv}(x \cup x')} = \sigma_x \vee \sigma_{x'}$ is in C. *Q.E.D.*

Recall that V is unique up to affine transformations, so we can normalize V by setting $V(\{1/K,\ldots,1/K\}) = 0$ and $V(x_{\sigma^c}) = c$. Now let $W: C \to \mathbf{R}$ be defined by $W(\sigma) = V(x_\sigma)$.

LEMMA 10: 1. *W is linear on C, i.e., $W(\sigma+\lambda\sigma')=W(\sigma)+\lambda W(\sigma')$, if σ, σ', and $\sigma+\lambda\sigma'$ are all in C.*

2. *W is continuous on C with respect to the sup norm topology.*

PROOF: 1. That W satisfies affinity, i.e., $W(\lambda\sigma+(1-\lambda)\sigma')=\lambda W(\sigma)+(1-\lambda)W(\sigma')$ follows immediately from Lemma 4. Our choice of normalization implies that W is linear: $W(\lambda\sigma)=W(\lambda\sigma+(1-\lambda)0)=\lambda W(\sigma)+(1-\lambda)W(0)=\lambda W(\sigma)$. Finally, then,

$$W(\sigma+\sigma')=2W\left(\frac{1}{2}\sigma+\frac{1}{2}\sigma'\right)=W(\sigma)+W(\sigma').$$

(Usually we would not think of V as linear in this sense, since $x\in X\Rightarrow \lambda x\notin X$. But if we "define" $\{1/K,\ldots,1/K\}$ as 0 and so define λx to be $\lambda x+(1-\lambda)\{1/K,\ldots,1/K\}$, then V is linear in this fashion as well.)

2. This follows from continuity of V and Lemma 8. *Q.E.D.*

In this part of the proof, we extend W to $C(S^K)$ in a series of steps. First, we restrict W to $C_+\equiv\{\sigma\in C\mid \sigma(s)\geq 0 \text{ for all } s\}$. Note that all the properties of C described in Lemma 9 hold for C_+. Next, define rC_+ to be the set of functions equal to r times some function in C_+ and let $H=\bigcup_{r\geq 0} rC_+$. Finally, let

$$H^*=H-H=\{\sigma\in C(S^K)\mid \sigma=\sigma^1-\sigma^2, \text{ for some } \sigma^1,\sigma^2\in H\}.$$

Now extend W to H^* by linearity. Specifically, for any $\sigma\in H$, there is an r such that $(1/r)\sigma\in C_+$, so define $W(\sigma)=rW((1/r)\sigma)$. Similarly, for any $\sigma\in H^*$, there are σ^1 and σ^2 such that $\sigma^i\in H$, for $i=1,2$, so let $W(\sigma)=W(\sigma^1)-W(\sigma^2)$. That these definitions do not depend on the precise r and σ^i chosen follows from the linearity of W (see Lemma 10). To extend W to $C(S^K)$, we show that H^* is dense in $C(S^K)$, so we can extend W by continuity since all points in $C(S^K)$ that are not in H^* are limits of points in H^*.

LEMMA 11: *H^* is dense in $C(S^K)$.*

PROOF OF LEMMA: By the Stone-Weierstrass theorem (see, e.g., Meyer-Nieberg (1991, Theorem 2.1.1, page 51)), we only need to show that: 1. H^* is a vector sublattice of $C(S^K)$; 2. H^* separates the points of S^K; 3. H^* contains the constant function 1_{S^K}.

Step 1: First note that H is a convex cone (i.e., a convex set that is closed under positive scalar multiplication), since it equals $\bigcup_{r>0} rC_+$ and C_+ is convex and contains the zero function.[32] Lemma 9 implies that H contains the supremum of any two of its elements. Next, note that $C(S^K)$ is a vector lattice, i.e., an ordered vector space that is a lattice (that is, contains the supremum and infimum for any two elements of $C(S^K)$).[33]

[32] The details are as follows. If $f\in H$, and $t\in R_+$ then clearly $tf\in H$. If $f_i\in H$, $i=1,2$, then $f_i=r_ig_i$, $g_i\in C_+$, $i=1,2$; say $r_2\leq r_1$. So if $\lambda\in[0,1]$ then

$$\lambda f_1+(1-\lambda)f_2=\lambda r_1g_1+(1-\lambda)\frac{r_2}{r_1}r_1g_2=r_1\left[\lambda g_1+(1-\lambda)\frac{r_2}{r_1}g_2\right].$$

But the function in square brackets is in C_+ because $(r_2/r_1)g_2\in C_+$. This last statement follows in turn because $(r_2/r_1)\in(0,1)$, and C_+ is a convex set that contains the zero.

[33] We defined the supremum, $\sigma\vee\sigma'$, in part (4) of Lemma 9; the infinimum, denoted $\sigma\wedge\sigma'$, is defined similarly. For $\sigma,\sigma'\in C(S^K)$, and for $r\in\mathbf{R}$, addition, $\sigma+\sigma'$, and scalar multiplication, $r\sigma$, are both defined in the usual way, under which $C(S^K)$ is obviously a vector space. It is ordered in the usual way and for that order it is an ordered vector space. Moreover, it also obviously contains the sup and inf of any two of its elements.

Now we show that since $H^* = H - H$, where H is a convex cone that includes the supremum of its elements, and H^* is a subset of a vector lattice, we can conclude that H^* is a vector sublattice. That H^* is an ordered vector space is trivial. That it includes the supremum of any two of its elements follows from the fact that H does. To see this, first note that $(\sigma_1 - \sigma_2) \vee (\sigma_1' - \sigma_2') = [(\sigma_1 + \sigma_2') \vee (\sigma_1' + \sigma_2)] - (\sigma_2 + \sigma_2')$, because $(\sigma_1 - \sigma_2) \vee \sigma = (\sigma_1 \vee \sigma + \sigma_2) - \sigma_2$ for any $\sigma \in H^*$. Using this we prove that $(\sigma_1 - \sigma_2) \vee (\sigma_1' - \sigma_2') \in H^*$. The elements $\sigma_1 + \sigma_2'$, $\sigma_1' + \sigma_2$, and $\sigma_2 + \sigma_2'$ are all in H; therefore $(\sigma_1 + \sigma_2') \vee (\sigma_1' + \sigma_2) \in H$ since it is closed under taking supremums. Therefore $(\sigma_1 - \sigma_2) \vee (\sigma_1' - \sigma_2') \in H^*$ from the preceding argument and the definition of H^*. Finally we prove that it includes the infimum of two of its elements. While H is not closed under taking infimums, this follows for $H^* = H - H$ by taking negatives. Specifically,

$$\begin{aligned}(\sigma_1 - \sigma_2) \wedge (\sigma_1' - \sigma_2') &= -[(\sigma_2 - \sigma_1) \vee (\sigma_2' - \sigma_1')] \\ &= -[((\sigma_1 + \sigma_2') \vee (\sigma_1' + \sigma_2)) - (\sigma_2 + \sigma_2')] \\ &= (\sigma_2 + \sigma_2') - [(\sigma_1 + \sigma_2') \vee (\sigma_1' + \sigma_2)].\end{aligned}$$

Now repeat the preceding argument. *Q.E.D.*

Step 2: Let $s, s' \in S^K$, $s \neq s'$. Note first that for any $x \in X$ which contains $(1/K, \ldots, 1/K)$, one has $\sigma_x \in C_+$. Now it is easy to construct a set with this property such that $\sigma_x(s) > \sigma_x(s')$. Find an element $\alpha \in \mathbf{R}^B$ such that $(s, \alpha) > \max\{0, (s', \alpha)\}$ (where (s, α) is the inner product—this can be done, for instance, by appeal to the separation theorem), and $\sum_{i \in B} \alpha^i = 1$. For λ small enough, $\lambda\alpha + (1 - \lambda)(1/K, \ldots, 1/K) \equiv \alpha(\lambda) \in \Delta(B)$, and if we let

$$x \equiv \{\theta\alpha(\lambda) + (1 - \theta)(1/K, \ldots, 1/K) \mid \theta \in [0, 1]\}$$

then we have $\sigma_x(s) = \lambda(s, \alpha) > \sigma_x(s')$ as claimed. *Q.E.D.*

Step 3 follows from Lemma 9 and the definition of H. *Q.E.D.*

LEMMA 12: *There exists a constant κ such that for all $f \in H^*$, $W(f) \leq \kappa \|f\|$ where $\|f\|$ is the supremum norm in $C(S^K)$.*

PROOF OF LEMMA: By compactness of X and the continuity of V, we know that there are best and worst sets in X. Let x denote a best set and y a worst set in X. By nontriviality, $x \succ y$. Let $\mathscr{B}_\varepsilon(\sigma_y)$ denote the subset of $C(S^K)$ within ε of σ_y. By continuity, there exists an $\varepsilon > 0$ such that for all $f \in H^* \cap \mathscr{B}_\varepsilon(\sigma_y)$, $W(f) < W(\sigma_x)$. Because H^* is closed under addition, this implies that for every $z \in H^*$ with $\|z\| < \varepsilon$, $W(z + \sigma_y) < W(\sigma_x)$. By linearity, $W(z) + W(\sigma_y) < W(\sigma_x)$ or $W(z) < W(\sigma_x) - W(\sigma_y)$. Equivalently, for all $z \in H^*$ with $\|z\| < 1$,

$$W(\varepsilon z) < W(\sigma_x) - W(\sigma_y)$$

or $W(z) < [W(\sigma_x) - W(\sigma_y)]/\varepsilon \equiv \kappa$. So for every $f \in H^*$,

$$W(f) = \|f\| W\left(\frac{f}{\|f\|}\right) < \|f\| \kappa. \qquad \textit{Q.E.D.}$$

LEMMA 13: *The functional W on H^* has a unique extension to a continuous linear functional on $C(S^K)$.*

PROOF OF LEMMA: H^* is a subspace of $C(S^K)$ and W is a real linear functional on H^*. Given the bound established by Lemma 12, we can apply the Hahn-Banach theorem (see Theorem 4, page 187 of Royden (1968)) to conclude W has an extension to a continuous linear functional on $C(S^K)$. That this extension is unique follows from the fact that H^* is dense in the supremum norm in $C(S^K)$, as shown in Lemma 11. *Q.E.D.*

By the Riesz representation theorem (see, e.g., Royden (1968)), every linear functional can be represented as integration against a measure. Hence V can be so represented.

PROPOSITION 3: *There is a measure μ on the Borel subsets of S^K such that for all $f \in C(S^K)$,*

$$W(f) = \int_{S^K} f(s)\mu(ds).$$

Thus, letting S be the support of μ on S^K, we have for all $x \in X$,

$$V(x) = \int_{S^K} \sigma_x(s)\mu(ds) = \int_S \max_{\beta \in x} U(\beta, s)\mu(ds).$$

This gives our additive EU representation. *Q.E.D.*

C.6. *Proof of Theorems* 1.C *and* 2

To show Theorem 1.C, recall that in the proof of Theorem 1.A, we showed that every closed sufficient set contains the subjective state space we constructed, the set of strongly relevant ex post preferences. Since the subjective state space for any weak EU representation must be sufficient, the subjective state space we constructed is contained in the closure of the subjective state space of any other weak EU representation. Because S^K contains all possible EU preferences, it is convenient to translate any alternative subjective state space to S^K. So let S^* denote the state space identified in our construction and let $S' \subseteq S^K$ be the set in S^K corresponding to the subjective state space for any other weak EU representation, so $S^* \subseteq \text{cl}(S')$. Suppose this containment is strict. Since $\text{cl}(S') \setminus S^*$ is open, there must be some $s \in S' \setminus S^*$. (If not, there is an open set in $\text{cl}(S')$ that does not intersect S', a contradiction.) By the definition of S^*, there is a neighborhood N of s such that for any $x, x' \in X$ with $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$, we must have $x \sim x'$. Since any subset of N will also have this property, we can assume that $N \subseteq S' \setminus S^*$. But then s is not relevant, contradicting its inclusion in S'.

Given this, it is easy to prove Theorem 2. Starting with Part 3, for ex ante preference $\succ_i$, let S_i^* be the set of s such that $\succ_s^*$ is strongly relevant for $\succ_i$. That is, it is the set of s such that for every neighborhood N of s, there are menus x and x' such that $x \nsim_i x'$ and $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$. Let $\mathbf{P}_i = \mathbf{P}(S_i^*, U_i)$.

The proof above of Theorem 1.C shows that $\mathbf{P}_i$ is the unique (closure of the) subjective state space for a weak EU representation of $\succ_i$. Hence what we need to show is that if $\succ_2$ is more uncertain than $\succ_1$, then $\mathbf{P}_1 \subseteq \mathbf{P}_2$, which is equivalent to showing $S_1^* \subseteq S_2^*$. So fix any $s \in S_1^*$. By definition, for every neighborhood N of s, there are menus x and x' with $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$ with $x \nsim_1 x'$. Because $x \nsim_1 x'$, it must be true that either $x \cup x' \nsim_1 x$ or $x \cup x' \nsim_1 x'$. Without loss of generality, suppose $x \nsim_1 x \cup x'$, so by Theorem 1.A and IR, $\text{conv}(x \cup x') \nsim_1 x'$. Note that for all $s' \in S^K \setminus N$,

$$\sigma_{\text{conv}(x \cup x')}(s') = \max\{\sigma_x(s'), \sigma_{x'}(s')\} = \sigma_x(s').$$

$\succ_2$ is more uncertain than $\succ_1$ implies that $x \nsim_2 x \cup x'$, so $x \nsim_2 \text{conv}(x \cup x')$. By the above, $\sigma_x(s') = \sigma_{\text{conv}(x \cup x')}(s')$ for all $s' \in S^K \setminus N$. Since N is arbitrary, $s \in S_2^*$.

The proof of Theorem 2 Parts 1 and 2 is similar. Fix a representation (S, U, u) and let $\mathscr{P}$ denote the closure of its set of positive states. First, we claim that the set of ex post preferences corresponding to S_+^* must equal $\mathscr{P}$. As before, it is convenient to think of S and $\mathscr{P}$ as subsets of S^K. Recall that $s \in S_+^*$ implies that for every open N, there are menus x and x' with $x \subset x'$, $x' \succ x$, and $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$. Hence $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S \setminus N$, implying $s \in \mathscr{P}$. Hence $S_+^* \subseteq \mathscr{P}$.

Suppose this inclusion is strict. Then there is an $s \in \mathscr{P} \setminus S_+^*$. Because S_+^* is closed, there is a neighborhood N of s such that $N \subset \mathscr{P} \setminus S_+^*$. By definition of $\mathscr{P}$, there must be x and x' with $x \subset x'$, $x' \succ x$, and $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S \setminus N$. By continuity of expected utility, then, the same is true for all $s' \in \text{cl}(S) \setminus N$. However, we know that $S_+^* \subseteq S^*$ and that S^* must equal the closure of S, so

$S^*_+ \subseteq \mathrm{cl}(S) \setminus N$. Hence we have x and x' with $x \subset x'$, $x' \succ x$, and $D(x, x') \cap S^*_+ = \varnothing$, contradicting Lemma 6. Hence $\mathscr{P} = S^*_+$. The analogous argument shows that $\mathscr{N} = S^*_-$.

So fix $\succ_1$ and $\succ_2$, where $\succ_2$ desires more flexibility than $\succ_1$, and any $s \in \mathscr{P}_1$. The above implies that for every neighborhood N of s (in S^K), there are menus x and x' with $x \subset x'$, $x' \succ_1 x$, and $\sigma_x(s') = \sigma_{x'}(s')$ for all $s' \in S^K \setminus N$. Since $\succ_2$ desires more flexibility, $x' \succ_2 x$. Since this is true for every neighborhood N, we must have $s \in \mathscr{P}_2$. Hence $\mathscr{P}_1 \subseteq \mathscr{P}_2$. An analogous argument covers the negative states. *Q.E.D.*

C.7. *Proof of Theorem* 3, *Parts B and C*

This proof makes use of a proposition that is an adaptation and generalization of Kreps' (1979) Theorem 2. This proposition, in turn, makes use of the following lemma. Given an ordinal representation, $R = (S, U, u)$, let $\mathbf{P}_R = \mathbf{P}(S, U)$. For any preference $\succ^*$ over $\Delta(B)$ and any $\beta \in \Delta(B)$, let

$$\mathscr{L}_{\succ^*}(\beta) = \{\beta' \in \Delta(B) \mid \beta \succsim^* \beta'\}.$$

That is, $\mathscr{L}_{\succ^*}(\beta)$ is the (weak) lower contour set for $\succ^*$ at β. Let $\mathscr{L}_{\succ^*}$ denote the collection of these lower contour sets for $\succ^*$ and let

$$\mathscr{L}_R = \bigcup_{\succ^* \in \mathbf{P}_R} \mathscr{L}_{\succ^*}.$$

LEMMA 14: *For any ordinal representation R of* $\succ$ *and any* $x \in \mathscr{L}_R$, $x \prec x \cup \{\beta\}$ *for any* $\beta \notin x$. *Also,* x *is convex if* $\succ$ *satisfies IR.*

The proof is straightforward and so is omitted.

PROPOSITION 4: *Let* $R_i = (S_i, U_i, u_i)$, $i = 1, 2$, *denote ordinal representations of an ex ante preference* $\succ$. *Then for every ex post preference in the state space of* R_1 *and every lower contour set* x *for that preference,* x *equals the intersection of some collection of lower contour sets in* R_2. *More precisely, for all* $\succ^* \in \mathbf{P}_{R_1}$, *for all* $x \in \mathscr{L}_{\succ^*}$, *there is an index set* N, *a nonrepeating sequence of states in representation* 2, $\{s_n\}_{n \in N} \subseteq S_2$, *and a sequence of lower contour sets* $\{x_n\}_{n \in N}$ *such that* $x_n \in \mathscr{L}_{\succ^*_{s_n}}$ *and*

$$x = \bigcap_{n \in N} x_n.$$

PROOF OF PROPOSITION: Let x denote any element of $\mathscr{L}_{R_1}$. For each $s \in S_2$, let

$$x_s = \left\{ \beta' \in \Delta(B) \;\middle|\; U_2(\beta', s) \leq \sup_{\beta \in x} U_2(\beta, s) \right\}.$$

Clearly, each $x_s \in \mathscr{L}_{R_2}$. Let

$$x' = \bigcap_{s \in S_2} x_s.$$

Clearly, $x \subseteq x_s$ for all $s \in S_2$, so $x \subseteq x'$. We now show that $x = x'$.

The proof that $x = x'$ is by contradiction, so suppose that x is a strict subset of x'. Let $\beta' \in x' \setminus x$. Clearly, $x \subseteq x \cup \{\beta'\} \subseteq x'$. Because these are ordinal representations, $\succ$ must satisfy monotonicity. Hence $x \precsim x \cup \{\beta'\} \precsim x'$. However, by the construction of x', for every $s \in S_2$,

$$\sup_{\hat{\beta} \in x'} U_2(\hat{\beta}, s) \leq \sup_{\hat{\beta} \in x} U_2(\hat{\beta}, s),$$

so the fact that the aggregator is increasing implies $x' \precsim x$. Hence $x \sim x \cup \{\beta'\}$. But, by Lemma 14, $x \in \mathscr{L}_{R_1}$ implies $x \cup \{\beta'\} \succ x$, a contradiction. Hence $x = x'$. *Q.E.D.*

We now prove Theorem 3.B and C.

3.B: Fix an ordinal EU representation with finite subjective state space $\mathbf{P}$. Let (S^o, U^o, u^o) be any ordinal representation of the same ex ante preference and let $\mathbf{P}^o$ be its subjective state space. Clearly, the result holds if $\mathbf{P}^o$ is infinite, so assume it is also finite.

We construct a function $f: \mathbf{P} \times \text{int}(\Delta(B)) \to \mathbf{P}^o$ (where $\text{int}(\Delta(B))$ is the interior of $\Delta(B)$). Fix any $\succ^* \in \mathbf{P}$ and any $\beta \in \text{int}(\Delta)(B)$. Because β is interior and $\mathbf{P}$ is a collection of EU preferences, $\mathscr{L}_{\succ^*_1}(\beta) \not\subseteq \mathscr{L}_{\succ^*_2}(\beta)$ whenever $\succ^*_1 \neq \succ^*_2$. Let $x = \mathscr{L}_{\succ^*}(\beta)$.

From the fact that all the preferences in $\mathbf{P}$ are expected-utility preferences, we see that x is a half-space. By Proposition 4, x must be the intersection of some collection of lower contour sets in $\mathscr{L}_2 \equiv \bigcup_{\succ^* \in \mathbf{P}^o} \mathscr{L}_{\succ^*}$. One way this can happen is if x itself is in $\mathscr{L}_2$. If so, let $f(\succ^*, \beta)$ equal the preference in $\mathbf{P}^o$ that generates this lower contour set.

So suppose $x \notin \mathscr{L}_2$. Then there exists an index set N, a nonrepeating sequence of states $s_n \in S^o$ for $n \in N$, and lower contour sets x_n for ex post preference $\succ^*_{s_n}$ for $n \in N$ such that

$$\bigcap_{k \in K} x_k = \mathscr{L}_{\succ^*}(\beta).$$

Without loss of generality, we can assume that if s and s' are distinct states in this sequence, then $\succ^*_s \neq \succ^*_{s'}$ since effectively only the smaller lower contour set appears in the intersection. Hence N must be smaller than the cardinality of $\mathbf{P}^o$ and so is finite. Because $\succ$ has an ordinal EU representation, Theorem 1.A implies that it must satisfy IR. Hence Lemma 14 implies that each x_n must be convex. It is impossible for a finite intersection of convex sets to equal a half-plane unless one of the sets is the half-plane, so there is some $\succ^*_o \in \mathbf{P}^o$ such that $\mathscr{L}_{\succ^*}(\beta) \in \mathscr{L}_{\succ^*_o}$. As noted, $f(\succ^*, \beta)$ is any such $\succ^*_o$.

We claim that $f(\cdot, \beta)$ is one-to-one. To see this, recall that for any $\beta \in \text{int}(\Delta(B))$, none of the $\mathscr{L}_{\succ^*}(\beta)$ sets is contained in any other. Hence there is no ex post preference relation, expected-utility or otherwise, which has more than one of these as a lower contour set. Also, for every $\succ^* \in \mathbf{P}$, $\mathscr{L}_{\succ^*}(\beta)$ is a lower contour set for $f(\succ^*, \beta)$. Hence $f(\succ^*_1, \beta) = f(\succ^*_2, \beta)$ iff $\succ^*_1 = \succ^*_2$, so $f(\cdot, \beta)$ is one-to-one. Hence $|\mathbf{P}^o| \geq |\mathbf{P}|$.

We now show by contradiction that $\mathbf{P}^o = \mathbf{P}$ or else $|\mathbf{P}^o| > |\mathbf{P}|$. So suppose $\mathbf{P}^o \neq \mathbf{P}$ but that $|\mathbf{P}^o| = |\mathbf{P}|$. Note that if $f(\succ^*, \beta) = \tilde{f}(\succ^*)$ for every β, then $\succ^*$ has the same lower contour sets as $\tilde{f}(\succ^*)$, implying $\succ^* = \tilde{f}(\succ^*)$. Hence $\mathbf{P}^o \neq \mathbf{P}$ implies that there is some $\succ^*_1 \in \mathbf{P}$ and some β^1, β^2 in the interior of $\Delta(B)$ such that $f(\succ^*_1, \beta^1) \neq f(\succ^*_1, \beta^2)$. Letting $\succ^*_o = f(\succ^*_1, \beta^1)$, the fact that $|\mathbf{P}^o| = |\mathbf{P}|$ implies that there must be some $\succ^*_2 \neq \succ^*_1$ such that $\succ^*_o = f(\succ^*_2, \beta^2)$. Furthermore, for every β', there must be some $\succ^* \in \mathbf{P}$ such that $\mathscr{L}_{\succ^*}(\beta') \in \mathscr{L}_{\succ^*_o}$. We now show that this cannot occur.

To see this, note that the ordinal EU representation has parallel indifference curves for each of its ex post preferences so the indifference curve for ex post preference $\succ^*_o$ through β^1 has a different slope than the same preference's indifference curve through β^2. Consider the line between β^1 and β^2. For each point β on this line, the indifference curve through β for preference $\succ^*_o$ must be the indifference curve through β for some preference $\succ^* \in \mathbf{P}$. Since $\mathbf{P}$ is finite, the indifference curve must have one of finitely many slopes. For any such β, then, there is a $\varepsilon_\beta > 0$ such that for every β' that is a distance less than ε_β away from β, the slope of the $\succ^*_o$ indifference curve through β is the same as the slope of the $\succ^*_o$ indifference curve through β' to ensure that distinct indifference curves never intersect. Let d be the infimum of the distance from β^1 along this line to a point β' where the slope of the $\succ^*_o$ indifference curve through β' differs from the slope at β^1. Since β^2 is such a point, d must be weakly less than the distance to β^2. By the argument above, we see that $d > 0$. Let β' be the point a distance d along the line from β^1. Suppose the slope of the $\succ^*_o$ indifference curve through β' differs from the slope through β^1. Then moving an arbitrarily small distance back toward β^1 must reach a point β'' where the slope of the indifference curve through β'' differs from the slope through β', a contradiction. So the slope at β' must be the same as that through β^1. But then moving an arbitrarily small distance from β' toward β^2 must reach a point β'' where the slope of the indifference curve through β'' differs from that through β', again a contradiction. Hence the slope through β^1 equals the slope through β^2, a contradiction.

3.C: The proof of part B shows that if there is an ordinal representation with a finite subjective state space, then the subjective state space for any ordinal EU representation must be of smaller cardinality. Hence if there is no ordinal EU representation with a finite subjective state space, there cannot be any ordinal representation with a finite subjective state space.

To show that an ordinal EU representation with a countable subjective state space must exist, suppose (S_1, U_1, u_1) is an ordinal EU representation of $\succ$ with uncountable subjective state space $\mathbf{P}_1$. Let $\mathbf{P}_2$ denote a countable subset of $\mathbf{P}_1$ that is dense in $\text{cl}(\mathbf{P}_1)$ and let S_2 denote the associated countable collection of states. Let U_2 denote U_1 restricted to S_2. Finally, define u_2 on $\mathscr{U}_2^*$ as follows. Given $U_2^* \in \mathscr{U}_2^*$, compute $U_1^*(U_2^*) \in \mathscr{U}_1^*$ by

$$U_1^*(U_2^*)(s) = U_2^*(s) \quad \text{if} \quad s \in S_2$$

and

$$U_1^*(U_2^*)(s) = \lim_{n \to \infty} U_2^*(s_n')$$

for some sequence $s_n' \to s$ (where this refers to convergence of $\succ_{s_n'}^*$ to $\succ_s^*$). By the denseness of $\mathbf{P}_2$, such a sequence must exist. Then define $u_2(U_2^*) = u_1(U_1^*(U_2^*))$. It is easy to see that for any set x, $u_2(u_2^*(x)) = u_1(U_1^*(x))$, so this generates an ordinal EU representation of the same ex ante preference that has a countable state space. *Q.E.D.*

C.8. *Proof of Theorem* 4.*B*

Let (S_1, U_1, u_1) denote any weak EU representation of $\succ$ and let (S_2, U_2, u_2) be the additive EU representation that Theorem 4.A tells us must exist. By definition, for any $U_2^* \in \mathscr{U}^*(S_2, U_2)$,

$$u_2(U_2^*) = \int_{S_2} U_2^*(s_2)\mu_2(ds_2).$$

By Theorem 1.B (see Definition 6 in Appendix B), this implies that, up to a monotone transformation, we have

$$u_1(U_1^*) = \int_{S_2} g(U_1^*)\mu_2(ds_2) = \int_{S_2} [\gamma(s_2)U_1^*(\pi(s_2)) + \delta(s_2)]\,\mu_2(ds_2)$$

for all $U_1^* \in \mathscr{U}_1^*$. Dropping out $\int_{S_2}\delta(s_2)\mu(ds_2)$ constitutes another monotonic transformation, so up to a monotone transformation,

$$u_1(U_1^*) = \int_{S_2} U_1^*(\pi(s_2))\hat{\mu}_2(ds_2)$$

where $\hat{\mu}_2(ds_2) = \gamma(s_2)\mu_2(ds_2)$. Now take the change of variables $s_1 = \pi(s_2)$ and let μ_1 be the resulting measure. Then we obtain

$$u_1(U_1^*) = \int_{S_1} U_1^*(s_1)\mu_1(ds_1),$$

completing the proof. *Q.E.D.*

REFERENCES

AL-NAJJAR, N., R. CASADESUS-MASANELL, AND E. OZDENOREN (1999): "Subjective Representation of Complexity," Northwestern University Working Paper.

ANSCOMBE, F. J., AND ROBERT J. AUMANN (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199–205.

CASTAING, CHARLES, AND M. VALADIER (1977): *Convex Analysis and Measurable Multifunctions*, Vol. 580 of *Lecture Notes in Mathematics*. New York, NY: Springer Verlag.

CLARK, FRANK H. (1983): *Optimization and Nonsmooth Analysis*, Canadian Mathematical Society Series. New York, NY: Wiley.

DEKEL, EDDIE, BARTON L. LIPMAN, AND ALDO RUSTICHINI (1998): "Recent Developments in Modeling Unforeseen Contingencies," *European Economic Review*, 42, 523–542.

DHILLON, AMRITA, AND JEAN-FRANÇOIS MERTENS (1996): "Relative Utilitarianism: An Improved Axiomatization," CORE Discussion Paper 9655.

DUNFORD, NELSON, AND JACOB T. SCHWARTZ (1958): *Linear Operators, Part I: General Theory*. New York, NY: Wiley.

ELLSBERG, DANIEL (1962): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643–669.

FISHBURN, PETER C. (1970): *Utility Theory for Decision Making*, Publications in Operations Research, No. 18. New York: Wiley.

GHIRARDATO, PAOLO (1996): "Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty," Mimeo, Cal Tech.

GUL, FARUK, AND WOLFGANG PESENDORFER (1999): "Temptation and Self-Control," Working Paper, Princeton University.

KELLEY, JOHN L. (1955): *General Topology*. New York: Springer-Verlag.

KEYNES, J. M. (1921): *A Treatise on Probability*. London: Macmillan.

KNIGHT, F. H. (1921): *Risks, Uncertainty, and Profit*. Boston: Houghton Mifflin.

KLEIN, ERWIN, AND ANTHONY C. THOMPSON (1984): *Theory of Correspondences: Including Application to Mathematical Economics*, Canadian Mathematical Society Series. New York, NY: Wiley.

KREPS, DAVID M. (1979): "A Representation Theorem for 'Preference for Flexibility'," *Econometrica*, 47, 565–576.

——— (1988): *Notes on the Theory of Choice*. Boulder, CO: Westview Press.

——— (1992): "Static Choice and Unforeseen Contingencies," in *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, ed. by Partha Dasgupta, Douglas Gale, Oliver Hart, and Eric Maskin. Cambridge, MA: MIT Press, 259–281.

MACLEOD, W. BENTLEY (1996): "Decision, Contract, and Emotion: Some Economics for a Complex and Confusing World," *Canadian Journal of Economics*, 29, 788–810.

MASKIN, ERIC, AND JEAN TIROLE (1999): "Unforeseen Contingencies and Incomplete Contracts," *Review of Economic Studies*, 66, 83–114.

MEYER-NIEBERG, PETER (1991): *Banach Lattices*. New York, NY: Springer Verlag.

NEHRING, KLAUS (1999): "Preference for Flexibility in a Savage Framework," *Econometrica*, 67, 101–120.

OZDENOREN, EMRE (1999): "Preference for Flexibility in an Anscombe-Aumann Framework," Northwestern University Working Paper.

ROCKAFELLAR, R. TYRELL (1972): *Convex Analysis*. Princeton, NJ: Princeton University Press.

ROYDEN, H. L. (1968): *Real Analysis*. New York, NY: Macmillan Publishing Co.

SAVAGE, LEONARD J. (1954): *The Foundations of Statistics*. New York, NY: Wiley.

SKIADAS, COSTIS (1997): "Subjective Probability under Additive Aggregation of Conditional Preferences," *Journal of Economic Theory*, 76, 242–271.