Discussion Paper No. 1202

A Unique Subjective State Space for Unforeseen Contingencies

by
Eddie Dekel
Barton L. Lipman
and
Aldo Rustichini

November 1997

A Unique Subjective State Space for Unforeseen Contingencies [1]

Eddie Dekel[2] Barton L. Lipman[3] Aldo Rustichini[4]

First Draft
November 1997

[1]We thank Peter Klibanoff, George Mailath, Jean François Mertens, Klaus Nehring, Phil Reny, Shmyoung Shi, Rani Spiegler, and numerous seminar audiences 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.

[2]Economics Dept., Northwestern University. E-mail: dekel@nwu.edu.

[3]University of Western Ontario. E-mail: blipman@julian.uwo.ca.

[4]CORE and CentER. E-mail: aldo@kub.nl.

Abstract

We axiomatically characterize a representation of preferences over opportunity sets which exhibit a preference for flexibility, interpreted as a model of unforeseen contingencies. In this representation, the agent acts as if she had a coherent prior over a set of possible future preferences, each of which is an expected utility preference. We show that the state space is essentially unique given the restriction that all future preferences are expected-utility preferences and is minimal even without this restriction. Finally, because the state space is identified, the additivity across states is meaningful in the sense that all representations are intrinsically additive.

1 Introduction

We derive a representation of preferences under unforeseen contingencies which extends that of Kreps (1979, 1992) by enriching the choice set and consequently obtaining much tighter results. We begin this introduction with an overview which briefly describes our view of unforeseen contingencies, Kreps' approach and the problems with it which we address, and our results. In the rest of the introduction, we devote a subsection to a more detailed explanation of each of these three issues. After that, we address possible criticisms of this approach and review the relevant literature.

1.1 Overview

It is commonplace to observe that real contracts are not as precisely detailed as the optimal contracts in standard models. One frequently given reason for this discrepancy is the presence of *unforeseen contingencies*. That is, real agents do not have a clear enough view of the set of possible situations to write such detailed contracts. Our goal in this paper is to provide a way of representing the behavior (preferences) of an agent who knows that she faces unforeseen contingencies.

To derive such a representation, we must first identify some aspect of potentially observable preferences which would provide useful information about the agent's perception of unforeseen contingencies. As noted by Kreps (1992), one behavior that seems linked to unforeseen contingencies is the desire to retain flexibility by not committing oneself in advance. Intuitively, if an agent cannot foresee future possibilities precisely enough, she should wish to retain some flexibility to react to the unexpected. Kreps uses preferences regarding such flexibility to derive a representation of unforeseen contingencies which looks much like the standard expected-utility model. In this sense, the surprising conclusion is that the standard model (albeit with some modifications) is also a model of unforeseen contingencies.[1]

To be more precise, Kreps considers preferences over what we refer to as *menus*. A menu is simply a subset of the available options and a choice of such an object is interpreted as a commitment to choose "in the future" from this set. Under two mild assumptions, Kreps (1979, 1992) derives a representation of preferences over such menus,

[1]Thus, as Kreps (1992) notes, the standard model is richer than we might have thought, although, as we will see shortly, the model is not completely standard since preferences are state dependent. Nevertheless, to argue that the model is not rich enough one needs to point out choice behavior that is not explained by the model.

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 a menu x, knowing that at some unmodeled second stage, he will learn the state of the world, s, and thus learn his preferences as represented by $U(\cdot, s)$. He then chooses the best object from the menu x according to his 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. One important point is that S and the $U(\cdot, s)$ functions are part of the representation, not a primitive of the model. In other words, even though the model does not assume that the agent foresees all possible future circumstances, the conclusion is that the agent acts *as if* he had a coherent view of the unforeseen contingencies. This view is summarized by the set of future preferences that he considers possible; we refer to this set of future preferences, $\{U(\cdot, s) \mid s \in S\}$, as the *subjective state space*.

We see two difficulties with using this model. First, the state space is not meaningfully identified; in particular it is not unique. As we argue in Section 1.3, this causes significant problems in developing applications of the model. Second, because of this, the additivity across states in the representation above is artificial in the sense that the additivity in the model is not a restriction on the preferences. As Kreps (*op. cit.*, page 567) points out: "the 'additive representation' . . . has limited significance — the representation is basically ordinal in character." More precisely, the set of preferences that can be represented in the form of (1) is the same as the set of preferences that can be represented by

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

for some increasing aggregator u. Moreover, some aggregators are "intrinsically" non-additive in the sense that they are not monotone transformations of additive functions. For preferences with such an intrinsically nonadditive representation, we can change the state space in such a way as to achieve an additive representation in the form of (1) on the new state space. In this sense, the inability to pin down the state space carries over to an inability to pin down the functional form of the representation.

We enrich the choice space by allowing menus of lotteries, instead of considering only menus of deterministic options. We give axioms which are necessary and sufficient for the existence of what we call an *additive EU representation* — a representation which is additive across states as in (1) above and where $U(\cdot, s)$ is an expected utility function for every s. In this class of representations, the state space is identified. In particular, we show (subject to caveats explained below) that every additive EU representation has the same state space. If we drop the restriction to additive representations and consider ordinal representations as in (2), the same result continues to be true. If we drop the restriction that every $U(\cdot, s)$ be an expected–utility function, there is a natural reason to focus on this particular state space: it is the smallest possible state space for any representation, additive or not. Finally, we also show that additivity is a necessary property of the representation. That is, given any ordinal representation with EU preferences for every s, the aggregator must be a monotone transformation of an integral.

The next three subsections explain the points above in greater detail, beginning with the unforeseen contingencies interpretation of the model. The introduction concludes with a discussion of potential criticisms and related literature. Section 2 presents the representation theorem and the subsequent section discusses the uniqueness and minimality of the state space and the sense in which additivity is meaningful. Most of the proofs are in the Appendix.

1.2 Modeling Unforeseen Contingencies

Clearly, the first point to clarify is exactly what we mean by "unforeseen contingencies." To contrast the idea with the usual Savage (1954) framework, recall that Savage takes as given a state space and set of consequences. A state is interpreted as a *complete* (in a sense explained below) specification of all relevant facts about the world. A consequence is viewed as a sufficiently rich description of an outcome as to completely determine how "well off" the agent is, independently of the state and the action choice which led to this consequence. Both states and consequences are viewed as exogenous and "objective" in the sense that they could be identified by an outside observer without knowing the agent's preferences. This is important in that if states and consequences are not objective in this sense, we cannot test the theory, even in principle. The actions available to the agent are formally represented as *acts*, which are functions from states to consequences. This is the sense in which the state space is complete: given any state, the agent knows precisely the consequences of any action.

We think of unforeseen contingencies as meaning that the individual has too vague a view of the set of possibilities for these assumptions to hold. To clarify, call a *situation* as precise a description of a possible state of the world as the agent can come up with. Roughly, we would think of a situation as corresponding to an event in the state space, rather than a single state. For example, if a state is described by the values of two relevant variables, x and y, then we can think of the state space as a set of (x, y) pairs. If the agent only thinks of x and not y, then a situation would be a possible value of x only, and hence could be thought of as corresponding to the set of (x, y) pairs in the true state space with this particular value of x. In the interesting case, the agent recognizes that x is not the only relevant variable, but, of course, does not know what variables are missing. That is, she knows that the situations are not complete descriptions of the world. Because of this, the agent is unsure how well off she would be given any particular action and situation. As an implication, even the most precise contract she can write will necessarily be vague or incomplete in the sense that it cannot pin down her payoff as a function of the state of the world that occurs since she can only write contracts as a function of the one variable, x, she has thought of. 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.

Intuitively, the key to modeling such an agent is to find a representation of how she perceives the payoffs from each action. A natural approach is to think of each possible specification of a payoff function as a state and to construct a new state space which would be complete. Thus, if the set of possible actions is A, we could *construct* a (subjective) state space where each state specifies the consequences for each and every action in A.[2] 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 consequence of her choice. Therefore, the payoff relevant contingencies are "how does each choice make me feel?" With this new 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.

On the other hand, directly assuming such a state space seems problematic. As we discuss in more detail in Section 1.5.2, one cannot simply apply the results of Savage (1954) or Anscombe and Aumann (1963) to such a state space for several reasons. Moreover, such an approach would formally involve, in a sense we make precise, exactly what we do herein. Finally, at a more conceptual level, defining states this way would, in effect, replace one "complete state–space" assumption with another, perhaps more plausible but still questionable, "complete state space" assumption.

The key insight of Kreps (1992) is that, instead of the above, we can use potentially observable behavior of the agent to uncover her view of how well off she would be as a function of her action. As mentioned above and discussed in more detail below, Kreps shows that if preferences regarding flexibility are sufficiently "well behaved," then the agent indeed acts as if she had a complete state space describing how her second stage preferences depend on the situation and acted like a standard expected utility maximizer with respect to this uncertainty. In other words, instead of *assuming* that the individual constructs a complete state space as described above, we *derive* the conclusion that the agent behaves *as if* she had such a state space. Naturally, we find the derivation more

[2] To the best of our knowledge, Fishburn [1970] was the first to discuss this interpretation of states of the world.

appealing. Also, it is instructive as it shows what conditions on preferences or behavior are required for the conclusion.

1.3 Summary of Kreps

We now turn to a brief review of Kreps (1979), as reinterpreted by Kreps (1992) in terms of unforeseen contingencies. For concreteness, the reader might want to consider an agent who must plan a meal for a given night in the distant future. Naturally, if there are unforeseen contingencies that would affect the agent's preferred meal on the night in question, the agent would prefer choosing a menu of options now, leaving the exact selection to be determined later. Alternatively, one can think of this as a stylized model of the allocation of control rights in a contract. Unforeseen contingencies then would influence what kinds of rights the agent wants to have.

Let B denote the finite set of (deterministic) options — food items in the first example — and consider the agents preferences,[3] $\succeq$, over subsets of B — menus — which are denoted as $x \in X = 2^B \setminus \{\emptyset\}$. If the agent faces no unforeseen contingencies, then he must know what his preferences will be over B. Denoting these preferences by $\succeq^*$, we could derive preferences over X 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 preferences over menus will not value flexibility. That is, no preference over menus which is generated in this way can have both $\{b, b'\} \succ \{b\}$ *and* $\{b, b'\} \succ \{b'\}$. In this sense, it is the desire for flexibility which reveals the existence of unforeseen contingencies.

Kreps (1979) isolates two key properties of preferences over X: *monotonicity*, or

$$x \supseteq x' \Longrightarrow x \succeq x'.$$

and what we call *the union condition*, or

$$x \sim x \cup x' \Longrightarrow x \cup x'' \sim x \cup x' \cup x''.$$

He shows that these conditions are necessary and sufficient for representing preferences $\succ$ by a *subjective state space* S, *state-dependent utility functions* $U : B \times S \to \mathbf{R}$, and a strictly increasing *aggregator* $u : \mathbf{R}^S \to \mathbf{R}$, as in (2).

That is, preferences over menus that satisfy the two axioms correspond precisely to behavior that would arise from a person who has a complete state space. Such a person behaves as if she conceives of future states S, and such that in each state, s, she chooses the maximal (according to $U(\cdot, s)$) element in the available set and aggregates these

[3]Throughout the paper preference orders are binary relations that are transitive and complete.

payoffs across states according to u. The only relevant aspect of S is the set of "second stage" preferences it induced (through the various $U(\cdot, s)$ functions) and so we refer to S and this collection of preferences interchangeably as the state space.

This representation is hardly pinned down at all. In contrast, say, to subjective expected utility preferences as derived by Savage (1954) or Anscombe and Aumann (1963) — where the probabilities in the representation are unique and the utility functions defined up to an affine transformation — here the subjective state space is not unique, very little can be said about the state-dependent utility functions themselves, and there are clearly no probabilities. Even in contrast to state dependent expected utility preferences — where, like here, probabilities are not pinned down — the indeterminacy in (2) of the state space is troubling for several reasons.[4] First, it clearly causes difficulties in deriving subjective probabilities. It seems impossible to identify the agent's probability distribution on the state space without identifying the latter.[5] 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 state space. Finally, applications of the model would seem to require some notion of "unforeseen contingency aversion" or a measure of the "extent" of unforeseen contingencies (both of which would probably be based on the on the size of S, and, loosely, on the variance of $U(\cdot, s)$ across states). If we cannot identify the 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.[6]

A second issue is that we might want a more structured representation. There are two components of the representation which one might want to restrict: (i) the state dependent utility functions, $U(\cdot, s)$, and (ii) the aggregator u. While the preferences that we take as a primitive, $\succ$, are over menus that are chosen today, the *representation* specifies (albeit, as discussed above, indeterminably) the second-stage preferences, $U(\cdot, s)$, which determine tomorrow's selection from the menu. We might think that tomorrow's preferences satisfy some additional properties, which would restrict the set of representations. In the framework above, the set B has no structure, so there is little one can ask of these second-stage preferences. Hence (i) requires enriching B in some fashion. Concerning (ii), as noted earlier, while Kreps (1979) does provide an additive subjective state dependent representation of the form (1), the additivity is not very meaningful.

[4] 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 which preserve preferences. However, there does not seem to be any simple, direct, useful statement of this set of transformations.

[5] It may be true that probabilities on relevant events are identified, however.

[6] For example, in Kreps' framework, we can have preferences represented by a single state — intuitively, preferences with no unforeseen contingencies — which are also represented by a large and quite "variable" state space.

1.4 Our Results

Our primary objective in this paper is to obtain a more structured representation of such preferences, where additivity is meaningful and, most importantly, where the subjective state space is pinned down. To do so, we extend preferences to a richer choice set, in particular where B includes lotteries, which is of independent interest.

Extending the preferences to sets of lotteries is important for two reasons. First, it is unrealistic and overly restrictive to assume that menus are chosen in a way that the options are deterministic. 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. Having enriched the choice environment to include sets of lotteries, it is natural to look for a more structured representation where the second stage preferences have an expected utility representation.

To clarify further the extent to which we satisfy our objectives, we state our main results, albeit a little vaguely because we leave some details for later. Consider preferences $\succ$ over subsets of $\Delta(B)$,[7] where a menu is a subset of $\Delta(B)$. Denote a generic element of B by b, a generic element $\Delta(B)$ by β (where $\beta(b)$ is the probability that β assigns to b), and a typical subset of $\Delta(B)$ by x. We also let b denote the lottery that assigns probability 1 to the outcome b. The preferences satisfy Kreps' monotonicity axiom, a natural *independence* axiom, and a *continuity* axiom (formally given in the next section) if and only if they have an additive EU representation — that is, iff they can be represented with a subjective state-space S and state-dependent expected utility functions $U : \Delta(B) \times S \to \mathbf{R}$ (i.e., satisfying $U(\beta, s) = \sum_{b \in B} \beta(b) U(b, s)$) as follows:

$$V(x) = \int_S \sup_{\beta \in x} U(\beta, s) \mu(ds). \tag{3}$$

Moreover, within the class of such representations, the state space (regarded as the set of possible second-stage preferences) is unique.[8] Furthermore, consider the class of all *ordinal* representations, i.e., those that have the form

$$V(x) = u[(\sup_{\beta \in x} U^o(\beta, s))_{s \in S^o}] \tag{4}$$

[7] In some contexts, e.g., if B is a set of dishes of food, this construction is artificial: we do not usually expect that an order of fish will yield chicken with high probability. (Although most readers will probably have experienced enough errors in orders that they would also not assign such a switch zero probability.) On the other hand, 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. Also, 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.

[8] Subject to caveats regarding infinite state spaces; see Section 3 for details.

for some collection of S^o, $U^o : \Delta(B) \times S^o \to \mathbf{R}$, and $u : \mathbf{R}^{S^o} \to \mathbf{R}$, where U^o need not be a representation of expected–utility preferences and u need not be additive. We define an *ordinal EU representation* to be an ordinal representation where each U^o is an expected utility function. Our uniqueness result extends to these representations as well: every ordinal EU representation has the same state space. Because our independence axiom is not necessary for the existence of an ordinal EU representation, this result does not rely on independence. Finally, our state space S is minimal within the class of all ordinal representations.

Because we can identify the state space, our additivity is also meaningful. More specifically, additivity is uniquely identified by our axioms in the sense that the aggregator for an ordinal EU representation of preferences satisfying our axioms is a monotone transformation of an integral. In this sense, the aggregator must be essentially additive. As a result, additivity *is* a restriction on preferences. Not all preferences with an ordinal representation and not even all those with an ordinal EU representation have an additive EU representation. By contrast, as noted above, all preferences that have the ordinal representation derived by Kreps', (2), do have an additive representation, (1).

These results achieve our main objectives: we identify an appealing representation with a unique state space which is minimal among all representations and in which additivity is meaningful. At a more detailed level, there are some other contributions here. As noted, we find assumptions on the preferences over sets that correspond to the *derived* state–dependent preferences being expected utility preferences, i.e., each $U(\cdot, s)$ being linear in $\Delta(B)$. As one might expect, continuity and a form of the independence axiom is needed, but we think the precise form is a little unexpected. Another surprising result is that the union condition used by Kreps is an implication of these axioms. Moreover, by pinning down an additive state dependent EU representation, we open the door to the next stage of pinning down probabilities, either by enriching the framework further or by adopting an approach such as that in Karni (1993). Our results also should make applications of this approach easier since the identification of the state space makes it possible to relate the structure of the state space to intuitive properties of preferences. For example, Theorem 12 shows that if one preference values flexibility more (and hence is more "averse" to unforeseen contingencies), then it must have a larger state space for its additive EU representation.

There is a sense in which our work relates to Kreps (1979) in the same way state independent expected utility representations of preferences over acts, as derived by Savage (1954) and Anscombe and Aumann (1963), relate to state dependent representations. Naturally, any state–independent representation could be written in a state–dependent form; moreover, there is no sense in which the typical usage of the state–independent form is "right." However, it seems like a natural normalization, it gives a useful structure and a unique meaningful subjective probability, and it is not the case that any state–dependent representation can be written in a state–independent form. Similarly, our additive expected utility representation could be written in an ordinal form (see (4)) and we do not claim that our representation is "right." However, it seems like a natural normalization, it gives a potentially useful structure and a unique meaningful subjective state space, and it is not the case that any preferences that have a representation in the form of (4) can be written in our additive EU form.

1.5 Discussion of the Approach

In this subsection, we briefly discuss three issues. First, we comment further on the appropriateness of this approach as a model of unforeseen contingencies. Second, we discuss in more detail our reasons for preferring to derive a state space from preferences over flexibility instead of postulating such a state space directly. Finally, we discuss the relationship between this approach and the incomplete contracts literature.

1.5.1 Criticisms of this Approach

There is something unappealing about viewing the states as second–stage preferences: if the agent is not aware of the actual factors which determine how she feels from each choice, how is she to form subjective probabilities over which preferences she is likely to have in the future?[9] While this is an important problem, it seems similar to the criticism that can be raised against the subjective expected–utility models of Savage and Anscombe–Aumann: where does the prior there come from? The conclusion in those papers is only that rational choice is equivalent to behavior *as if* there was a prior. Similarly here, the point is that *rational* choice in the presence of unforeseen contingencies is equivalent to creating a subjective state space and maximizing in (almost) the usual way. That is, the standard model applies, although where the prior comes from is at least as unclear as before.

To develop this analogy further, one of the fascinating contributions of Savage and Anscombe–Aumann was to show that even when there are no objective probabilities over the state space, the standard von Neumann–Morgenstern objective expected–utility model is appropriate. That is, the use of the model in an application does not change — we only change the interpretation of the prior to a subjective one. Similarly, this approach to unforeseen contingencies shows that the same model (albeit with state–dependent preferences) is again appropriate, but this time we change the interpretation of the state space to be subjective.

[9] While there is (as yet) no representation with meaningful probabilities, the criticism remains valid if we interpret it as asking how the individual is to make comparisons across states.

On the other hand, the Ellsberg paradox makes the point that observed behavior *is* sensitive to risk as well as uncertainty, and therefore that we should treat the case of unknown probabilities differently from the case where objective probabilities are given. Similarly, if our model is missing anything due to the existence of unforeseen contingencies, one must first find an example of the effect on behavior. The mere fact that this model doesn't seem to describe the *process* of decision making doesn't imply that it fails to be an accurate model.

A related criticism of our approach is that the conclusion suggests that something is wrong: a model of decision making in the presence of unforeseen contingencies should be different from the usual additive EU model, and if we don't get that, then we are assuming too much rationality. Our view is that, to the contrary, one should first explore the implications of weakening only the assumption of completely foreseen contingencies. We think it is important to understand that imposing very minimal consistency properties on the decision maker (that she has preference ordering satisfying monotonicity, continuity, and independence) implies that the decision maker behaves as if she has a subjective state space and is an additive EU maximizer. Why should these (more or less) standard axioms be more suspect in the presence of unforeseen contingencies? Moreover, while our perspective is obviously normative, the result also makes a descriptive point: any model of decision making that is inherently different will have to violate at least one of the axioms, which in turn suggests that it might be difficult to justify why such a model will have an EU structure, even on foreseen parts of the state space.

1.5.2 Problems with the Direct Approach

As mentioned in Section 1.2, one could proceed quite differently by simply postulating a state space which is equal to the set of all possible second stage preference relations and continuing with a more traditional analysis. We believe that this more direct approach has several philosophical drawbacks and that the only version of this approach which would avoid these would be equivalent to the approach we take here. Before commenting on these drawbacks, first note that the analysis would *not* follow directly from Savage because of the state dependence of the preferences.

One of the philosophical problems we have with the direct approach, noted earlier, is that it requires us to directly assume that the agent has a sufficiently precise and coherent view of the set of future possibilities as to be represented by such a state space, an assumption which seems at odds with the goal of studying unforeseen contingencies. One of the appealing aspects of Kreps' approach is that we can derive the subjective state space using surprisingly weak conditions by examining the agent's revealed preferences.

A second problem is that this approach runs counter to the traditional view of the state space and consequences as objective. The reason this is a problem is that it means that we cannot construct the acts over which we would need to observe the agent's preferences, and hence cannot hope to observe his preferences. That is, the direct approach is inconsistent with the revealed preference approach to decision theory. To make the point more concretely, suppose for simplicity that there are two actions, a and b, and no lotteries. To construct a state space equal to the set of preferences, then, we would need two states, say s_1 and s_2, where a is strictly preferred to b in s_1 and where b is strictly preferred in s_2.[10] Suppose we formalize this by supposing that an action has "payoff" 1 in the state where it is preferred and 0 in the state where it is not preferred. To carry out the traditional Savage analysis, we would need to consider all functions from the set of states $\{s_1, s_2\}$ into our consequence space, $\{0, 1\}$. Such an act can be written as a pair, $(u_1, u_2) \in \{0, 1\} \times \{0, 1\}$, where u_i is the payoff in state i. It is easy to imagine finding out whether the agent prefers $(1, 0)$ to $(0, 1)$ since this corresponds in an obvious way to whether the agent prefers a to b. Similarly, we could imagine learning whether the agent strictly prefers $(1, 1)$ to $(1, 0)$ or $(0, 1)$ by asking whether he would strictly prefer getting his favorite choice in each state to committing himself to either a or b. In principle, we could implement $(1, 1)$ by simply allowing the agent to take his choice between a and b after the state is realized. Note that this is precisely equivalent to considering the agent's preferences over menus: $(1, 1)$ corresponds to $\{a, b\}$; $(1, 0)$ to $\{a\}$; and $(0, 1)$ to $\{b\}$.

But how can we observe the agent's ranking of $(0, 0)$ relative to the other choices?[11] This act corresponds to giving the agent his least favorite option in both states. If — as in Savage — we were to observe the state, we could easily construct this act and offer it, so there would be no problem with observing the agent's preferences. But the very nature of the state space makes such observation wildly implausible. Hence there is no obvious procedure by which we can implement this act. Clearly, we cannot induce the agent to reveal the state to us by promising to give him his least favorite outcome in each state!

In other words, it is hard to imagine what such an act could possibly correspond to in terms of a real, potentially observable option. The acts which do have such an interpretation would be those which are *incentive compatible*, i.e., those that could be implemented

[10] To make the point most simply, we ignore the possibility of a state in which the agent is indifferent between the two.

[11] The agent should *weakly* prefer $(0, 1)$ and $(1, 0)$ to $(0, 0)$, but distinguishing between indifference and strict preference is crucial. Intuitively, if $(0, 0)$ is indifferent to $(1, 0)$ but $(0, 1)$ is strictly preferred to either, then this tells us that the agent gives s_1 probability zero. In principle, this preference could be deduced from other information. For example, if we assume the sure–thing principle, then $(1, 0)$ is strictly preferred to $(0, 0)$ if and only if $(1, 1)$ is strictly preferred to $(0, 1)$. Since we observe the latter, the sure thing principle enables us to deduce the former. Our point is that a simple application of the direct approach requires comparisons which we cannot observe even in principle. In fact, this kind of deduction is the basis of what we achieve: our independence axiom allows us to extend the preferences from the set of incentive compatible acts to the set of all acts. We do not prove the result this way because our version of this direct approach is more complex than the indirect approach.

by the agent truthfully revealing her preferences and then being given the appropriate outcome. As noted above, such incentive compatible acts naturally correspond to subsets of $\Delta(B)$. (Any subset, say x, of $\Delta(B)$ could be thought of as the act where in any state s, the agent receives that $\beta \in x$ which is optimal according to his preferences in state s. This act is clearly incentive compatible in the sense defined above. Similarly, any incentive compatible act corresponds to a set in this fashion simply by taking the range of the act.) Acts which are not incentive compatible, like $(0, 0)$, require the agent to be committed in some fashion which is difficult to interpret, and do not correspond to menus.[12]

Summarizing, then, we believe that the most appealing version of the direct approach would restrict attention to incentive compatible acts. Not only do these acts naturally correspond to menus, it is also true that the only way to construct a physical version of such an act is by offering the agent her choice from some fixed set — that is, via a menu. Finally, we strongly believe in the revealed-preference approach to decision making, and this requires considering menus. Given all this, and the fact that the menu approach does not commit us to assuming the existence of a subjective state space, we feel that this approach is superior to the direct approach.

1.5.3 Incomplete Contracts

We commented earlier that in the presence of unforeseen contingencies, the individual will view all contracts she can write as incomplete because she knows that some relevant variables are omitted. Hence, one expects that the agent will prefer contracts which leave some flexibility. Kreps (1992) notes that such preferences for flexible contracts could come from any source of contractual incompleteness. In this sense, there may not be any behavioral differences between unforeseen contingencies and other reasons for incomplete contracts.

The incomplete contracts approach basically *assumes* a representation much like this one, where some kinds of uncertainty are interpreted as *non-contractible*, an assumption generating a value for flexibility and hence for residual control rights (see Hart (1995)). As in Kreps (1992), we *derive* the representation from a preference for flexibility. This might enable us to state a converse to the claim that incompleteness of contracts leads to a value for residual control rights, namely that if residual control rights have value, then the contracting problem can be represented with non-contractible uncertainty. Whether one wants to call such a representation an incomplete-contracts model or a model with expected-utility preferences where the state space is subjective (and hence non-contractible) depends presumably on one's upbringing.[13]

[12] Nehring (1996) makes a similar observation.

[13] There are however some important steps unaddressed in developing such a converse: one needs to

It is also worth noting that this interpretation of the incomplete contracts model allows one to address a natural criticism of interpreting that model as one allowing unforeseen contingencies. In the absence of the framework offered here and in Kreps (1992), it would seem very strange to interpret the non-contractible uncertainty as unforeseen contingencies. The problem is that it is not clear how the agents could foresee these states in their evaluation of expected utility if they cannot foresee them well enough to contract on them. The answer is that these states are subjective states and hence non contractible. Equivalently, the agents can use these states in their maximization problem because they correspond, not to the actual situations that determine the utility, but to the utility itself. In fact, as the proofs of our results demonstrate, the states *are* utility functions. As Maskin and Tirole (1997) point out, the ability to foresee utility consequences is itself quite strong, leading to conclusions quite distinct from those usually drawn in the incomplete contracts literature. While we do not intend to take either side in this debate, we do argue that the framework used in the incomplete contracts literature is naturally interpretable as one where the states are subjective and correspond to state-dependent preferences.

1.6 Related Literature

The literature on unforeseen contingencies can be divided into epistemic and decision theoretic approaches.[14] The former category consists of papers which take the knowledge of an agent who faces unforeseen contingencies as the key to the model rather than action choices.[15] By contrast, the decision theoretic approach derives a representation of the agent from preferences, where some aspects of this representation can be interpreted as statements about the agent's knowledge or beliefs.

Aside from the aforementioned work of Kreps, the only papers we know of that take decision-theoretic approaches to this subject are Ghirardato (1995), Skiadis (forthcoming), and Nehring (1996). All three share our view that with unforeseen contingencies, the agent only recognizes situations, not states, and hence does not know her exact preferences conditional on any given situation. Ghirardato models this by assuming that the agent associates a set of consequences, rather than a single consequence, with each situation. Thus he gives a generalization of subjective expected utility to acts which 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 con-

extend this model to multiple agents and one needs to have some degree of state independence in order to obtain meaningful probabilities.

[14]For a survey, see Dekel, Lipman, and Rustichini (1997).

[15]See Fagin and Halpern (1988), Geanakoplos (1989), Modica and Rustichini (1993, 1994), and Dekel, Lipman, and Rustichini (forthcoming).

sequences. Skiadis studies preferences over actions conditional on situations and derives a representation where the agent has a subjective utility for each action conditional on each situation. Intuitively, this represents the agent's "expectation" of the utility consequences over the unforeseen aspects of a situation. The Kreps approach is similar to Skiadis' approach in that both use the agent's preferences to identify the utility consequences of acts as a function of the state. Finally, while Nehring, like us, follows Kreps' approach, he allows for preferences over acts over menus. That is, like Ghirardato, his acts are functions from states to sets, but unlike Ghirardato and like us, his representation involves an implicit second–stage at which the agent chooses from the appropriate set. Instead of an expected–utility restriction on the second–stage preferences, Nehring restricts attention to second stage preferences with the unappealing form in which there are only two (thick) indifference curves — that is, his second–stage preferences are represented by a utility function that takes only two values. He does give a uniqueness result given this restriction and, as a consequence, his additivity is meaningful in the same sense as ours. He does not have a minimality result — none is possible with his state space since it is typically not minimal.

2 The Additive Expected–Utility Representation

Let B be a finite set of n "prizes," let $\Delta(B)$ denote the set of probability distributions on B, and let $\mathcal{P}(\Delta(B))$ denote the set of nonempty subsets of $\Delta(B)$. A typical element of $\mathcal{P}(\Delta(B))$ will be denoted x (or $\tilde{x}$, x', etc.), while a typical element of $\Delta(B)$ will be denoted by β, β', or β''. We use λ for convex combinations. Suppose we have a preference relation $\succ$ on $\mathcal{P}(\Delta(B))$ which is a weak order. We consider an environment where the individual first chooses a set and at a later stage will choose among the elements of this set, but we do not explicitly model this second–stage choice. (In fact, as we discussed in the introduction, our representation does have implications for this second stage choice.) We impose three axioms on the preferences: monotonicity, continuity, and independence. The first is the same as Kreps' monotonicity, the second is essentially standard, and the third is our adaptation of the usual independence axiom.

Monotonicity

Monotonicity: $x \subseteq x' \implies x' \succeq x$.

Continuity

We endow $\mathcal{P}(\Delta(B))$ with the Hausdorff topology, τ_H; see the appendix for precise definitions. For every x, the *strict lower contour set* with respect to the preference order is defined as usual by $L(x) \equiv \{x' \subseteq \Delta(B) \mid x \succ x'\}$, and similarly the *strict upper contour set* is defined by $G(x) \equiv \{x' \subseteq \Delta(B) \mid x' \succ x\}$.

Continuity: The strict upper and lower contour sets are open. That is, for every $x \subseteq \Delta(B)$, $G(x) \in \tau_H$ and $L(x) \in \tau_H$.

Independence

To formally state the independence axiom, 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 $\lambda x + (1 - \lambda)x' \in \mathcal{P}(\Delta(B))$ 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 $\lambda\beta(b) + (1 - \lambda)\beta'(b)$.

Independence: For all $\lambda \in [0, 1]$, $x'' \in \mathcal{P}(\Delta(B))$,

$$x \succeq x' \Longrightarrow \lambda x + (1 - \lambda)x'' \succeq \lambda x' + (1 - \lambda)x''.$$

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 this condition by breaking it into two parts. Roughly speaking, our definition of convex combination is a kind of reduction of compound lotteries assumption. Given this, our independence axiom can be thought of as no different from the usual one.

To see this, first suppose we think of $\lambda x + (1 - \lambda)x''$ as a random determination of a set, giving the individual x with probability λ and x'' otherwise. (The justification for this identification is discussed next.) Then it is clear that this axiom is precisely the usual independence axiom and is interpreted in precisely the usual way: the difference between $\lambda x + (1 - \lambda)x''$ and $\lambda x' + (1 - \lambda)x''$ is only in the "λ" event, so the preference between these should be the same as the preference between x and x'.[16]

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. To see this most easily, suppose $x = \{\beta_1, \beta_2\}$ and $x'' = \{\beta_1'', \beta_2''\}$ and consider how the individual should

[16] Nehring (1996) considers preferences over lotteries over sets and uses precisely this form of the independence axiom. However, he does not follow our next step of identifying lotteries over sets with our definition of convex combinations of sets, an identification which is at the heart of our independence axiom.

view the gamble giving x with probability λ and x'' otherwise. The individual knows that whatever menu the gamble gives her at the first stage, she will choose her preferred element from that set at the second stage. Suppose, then, that she is considering the circumstances in which she would choose β from x and β'' from x'' at the second stage. In these circumstances, she would find this lottery over sets as equivalent to receiving the lottery $\lambda\beta + (1-\lambda)\beta''$ at the second stage. Since, in principle, her situation in the second stage could lead her to any particular pattern of choices from x and x'', she should be indifferent between this lottery over sets and receiving the set $\{\lambda\beta_1 + (1-\lambda)\beta_1'', \lambda\beta_2 + (1-\lambda)\beta_1'', \lambda\beta_1 + (1-\lambda)\beta_2'', \lambda\beta_2 + (1-\lambda)\beta_2''\}$ for sure. That is, she should view the gamble over sets as equivalent to the convex combination of sets as we defined it.

It is worth emphasizing that this is a normative argument, but that it does *not* depend on the agent being able to articulate the real circumstances that correspond to choosing β from x and β'' from x''. The argument above can be explained to an agent who knows that there are circumstances of which she is unaware and that she cannot describe precisely. All she needs be aware of is that she will eventually choose one of the elements in the set to which she has restricted herself at the first stage.

Let $\text{cl}(x)$ denote the closure of x (in the Euclidean topology on $\Delta(B)$).

Definition 1 *An* additive EU representation *of $\succ$ is a set S,*[17] *a (countably additive) measure μ with full support on S, and state-dependent utility functions $U : S \times \Delta(B) \to \mathbf{R}$ such that $V : \mathcal{P}(\Delta(B)) \to \mathbf{R}$ represents preferences, where*

$$V(x) = \int_S \sup_{\beta \in x} U(\beta, s)\, \mu(ds) = \int_S \max_{\beta \in \text{cl}(x)} U(\beta, s)\, \mu(ds)$$

and where each $U(\cdot, s)$ is an expected utility function in the sense that

$$U(\beta, s) = \sum_{b \in B} U(b, s)\beta(b).$$

The word "additive" refers to the additivity across S. Later, we consider representations where the payoffs across states are aggregated in a nonadditive fashion, as discussed in Section 1. The "EU" refers to the fact that each possible second stage preference is an expected utility preference.

Remark 1 When the state space is finite or countable, the measure μ is simply a normalization. That is, in either case, we can replace $U(\cdot, s)\mu(s)$ with a new utility function

[17] To be more precise, we require a measure space where U is measurable with respect to this space. Since we make no use of measurability considerations, we avoid the details.

$\bar{U}(\cdot, s)$ and eliminate the measure. In the uncountable case, the measure can be important since we must define integration with respect to some measure. It is easy to see that if μ is absolutely continuous with respect to Lebesgue measure, then we can replace it in a similar manner with Lebesgue measure. However, no obvious replacement is possible if μ is continuous but singular.

Theorem 1 *Preferences on $\mathcal{P}(\Delta(B))$ satisfy monotonicity, independence and continuity if and only they have an additive EU representation.*[18]

The necessity of the axioms for the representation is easily shown. The remainder of this section is devoted to the proof of sufficiency. Since this material is not necessary for an understanding of our subsequent results, a reader who is not interested in the proof can safely omit the remainder of this section. All proofs omitted below are contained in the appendix. We begin by establishing the existence of a representation of preferences that is linear on $\mathcal{P}(\Delta(B))$, Proposition 3 below. We then describe how we transform this into the desired representation.

Let $\text{conv}(x)$ denote the convex hull of x.

Lemma 2 *If preferences satisfy monotonicity and continuity, then for every $x \in \mathcal{P}(\Delta(B))$, $\text{cl}(x) \sim x$. If preferences also satisfy independence, then $\text{conv}(x) \sim x$.*

In light of this lemma, we henceforth restrict ourselves to the set of closed, convex, nonempty subsets of $\Delta(B)$, denoted by X.[19]

Proposition 3 *There is a linear $V : X \to \mathbf{R}$ that represents preferences, i.e.,*

$$x \succ x' \text{ iff } V(x) > V(x')$$

[18]Because Kreps' union condition ($x \sim x \cup x' \Rightarrow x \cup x'' \sim x \cup x' \cup x''$) is necessary for an additive EU representation, the theorem implies that it must be an implication of our axioms. A direct proof of this fact is not difficult. First, note that

$$\frac{1}{2}x + \frac{1}{2}[x \cup x' \cup x''] \subseteq \frac{1}{2}[x \cup x'] + \frac{1}{2}[x \cup x''],$$

so monotonicity implies that the right-hand side is weakly preferred. Suppose, then, that $x \sim x \cup x'$. By independence, then, we must have $x \cup x'' \succeq x \cup x' \cup x''$. By monotonicity, we get the oppposite weak preference, so this must be indifference, implying the union condition. We thank Klaus Nehring for showing us a critical step in this argument.

[19]The Hausdorff topology on X is metrizable and is equivalent to the Hausdorff topology on $\mathcal{P}(\Delta(B))$ restricted to X.

where

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

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

The proof of this proposition is straightforward. We verify that the usual mixture space axioms hold given our definition of convex combinations and then apply the Herstein and Milnor theorem (see, e.g., Fishburn (1970), Theorem 8.4, page 113, or Kreps (1988), page 54). The continuity statement is an immediate consequence of the uniqueness result and the existence of a continuous representation (see Lemma 15 in the appendix.) This is the function V we are going to explore in the rest of this section.

The rest of the proof is a little indirect, so we provide a rough sketch. We first identify each element x in X with a convex function σ_x defined on a subset of $\mathbf{R}^n$, denoted S^n.[20] (This set is formally defined below and the function is the support function — see Rockafellar (1972), page 28). These functions have the form of $\sigma_x(s) = \max_{\beta\in x} U(\beta, s)$ for $s \in S^n$, where U is linear in β, i.e., $U(\beta, s) = \sum_{b\in B} U(b, s)\beta(b)$. These will turn out to be the state dependent utility functions, and a subset of S^n will be the state space. Let C be the set of functions on S^n obtained from X; we show that the mapping from X to C is one to one. We then use V to define a monotonic, continuous, linear function, denoted W, on C. Since W is monotonic, continuous and linear, and due to the structure of C (details are provided below and in the appendix), W can be extended to a monotonic, continuous linear function on the set of all continuous functions on S^n. 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 $W(\sigma_x) = \int_{S^n} \sigma_x(s)\mu(ds)$, or $V(x) = \int_{S^n} \sigma_x(s)\mu(ds) = \int_{S^n} \max_{\beta\in x} U(\beta, s)\mu(ds)$, yielding the desired representation.

For convenience, we write $B = \{b_1, \ldots, b_n\}$. Let $S^n = \{s \in \mathbf{R}^n \mid \sum s_i = 0, \sum |s_i| = 1\}$, and let $C(S^n)$ denote the set of continuous functions on S^n. We order these functions pointwise as usual — that is, $\sigma \geq \sigma'$ means $\sigma(s) \geq \sigma'(s)$ for all $s \in S^n$. We now map X into $C(S^n)$, denoting the image of x by σ_x, where for any $s = (s_1, \ldots, s_n) \in S^n$,

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

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

$$x_\sigma \equiv \bigcap_{s\in S^n} \{\beta \in \Delta(B) \mid \sum_i \beta(b_i)s_i \leq \sigma(s)\}.$$

[20] Recall that n is the number of elements of B.

The following lemma implies that the mapping of X to C preserves the mixture space structure, i.e., preserves convex combinations; that monotonicity in terms of set inclusion corresponds to the natural monotonicity for functions; that x_σ and σ_x are inverses of one another; and that the Hausdorff topology on convex sets corresponds to the sup norm topology on the support functions of the sets .

Lemma 4 *For all $x, x' \in X$,*

1. $\sigma_{\lambda x+(1-\lambda)x'} = \lambda\sigma_x + (1-\lambda)\sigma_{x'}$
2. $x \subseteq x' \iff \sigma_x \leq \sigma_{x'}$
3. $x_{(\sigma_x)} = x$, *and* $\sigma_{(x_\sigma)} = \sigma$.
4. $d_{\text{Hausdorff}}(x, x') = d_{\text{supnorm}}(\sigma_x, \sigma_{x'})$.

Proof. These are standard results that follow immediately from the definitions; see, e.g., Clark (1983), Castaing and Valadier (1977), and Rockafellar (1972). ▌

The next lemma states some basic properties of C.

Lemma 5

1. *C is convex.*
2. *The zero function is in C, in particular $\sigma_{\{(1/n,...,1/n)\}}(s) = 0$ 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 $\sigma' \in C \Longrightarrow \sigma \vee \sigma' \in C$, where $(\sigma \vee \sigma')(s) = \max\{\sigma(s), \sigma'(s)\}$.*

Proof.

1. Given σ_x and $\sigma_{x'}$ in C, using Lemma 4, part (1), and the convexity of X, any convex combination of σ_x and $\sigma_{x'}$ is in C.
2. For any $s \in S^n$, we have $\sum_i s_i = 0$ so by definition $\sigma_{\{(1/n,...,1/n)\}}(s) = \sum_i \frac{1}{n} s_i = 0$.

3. First note that $\sigma_{\Delta(B)}(s) = \max_i\{s_i\} \geq 1/(2n)$. The equality follows from the definition of the support functions. The inequality follows from the definition of S^n. (If $\max_i\{s_i\} < 1/(2n)$, then $\sum_{\{i\ s_i>0\}} s_i < 1/2$. Then, since $\sum_i s_i = 0$, also $\sum_{\{i\ s_i<0\}} |s_i| < 1/2$. But then, $\sum_i |s_i| < 1$, which contradicts the definition of S^n.) Now consider $x \equiv \{\beta \mid \sum_i \beta(b_i)s_i \leq c$ for all $s \in S^n\}$, where $0 < c < 1/(2n)$. Clearly x is a closed, convex, and nonempty subset of $\mathbf{R}^n$. It is easy to see that we could have defined our mapping from X into $C(S^n)$ to have as its domain all convex, closed nonempty subsets of $\mathbf{R}^n$ without affecting Lemma 4. With this definition, clearly, σ_x is the constant function c. It remains to show that $x \in X$. By part (2) of Lemma 4, since $\sigma_x \leq \sigma_{\Delta(B)}$ we know that $x \subset \Delta(B)$, so $x \in X$.

4. Given σ_x and $\sigma_{x'}$ in C, it is easy to see that $\sigma_{\text{conv}(x,x')} = \sigma_x \wedge \sigma_{x'}$ is in C. ▌

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

Lemma 6

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.*

3. W *is monotonic with respect to the natural order on functions, i.e., if* $\sigma(s) \geq \sigma'(s)$ *for all* s, *then* $W(\sigma) \geq W(\sigma')$.

Proof.

1. That W satisfies "convexity," i.e., $W(\lambda\sigma + (1-\lambda)\sigma') = \lambda W(\sigma) + (1-\lambda)W(\sigma')$ follows immediately from Lemma 4, part (1). 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,

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

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

2. This follows from continuity of V and Lemma 4, part (4).

3. This follows from monotonicity of V with respect to set inclusion and Lemma 4, part (2). ▌

In this part of the proof, we extend W to $C(S^n)$. It turns out to be convenient to do this 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 5 hold for C_+. Next, define rC_+ to be the set of functions equal to r times some function in C_+ and let $H = \cup_{r>0} rC_+$. Finally, let

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

Now extend W to H^* by linearity. Specifically, for any $\sigma \in H$, there is an r such that $\frac{1}{r}\sigma \in C_+$, so define $W(\sigma) = rW(\frac{1}{r}\sigma)$. 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 6). To extend W to $C(S^n)$ we show that H^* is dense in $C(S^n)$, and then show that we can extend W by continuity since all points in $C(S^n)$ that are not in H^* are limits of points in H^*.

Lemma 7 *H^* is dense in $C(S^n)$.*

Lemma 8 *The functional W on H^* has a unique extension to a continuous and monotonic linear functional on $C(S^n)$.*

Now from 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 9 *There is a probability measure μ on the Borel subsets of S^n such that for all $f \in C(S^n)$,*

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

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

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

This proposition only characterizes $V(x)$ for closed and convex $x \subseteq \Delta(B)$. The following lemma extends the characterization to all nonempty subsets of $\Delta(B)$, completing the proof of Theorem 1.

Lemma 10 *For every $x \in \mathcal{P}(\Delta(B))$,*

$$V(x) = \int_S \sup_{\beta \in x} U(\beta, s)\, \mu(ds).$$

3 Uniqueness Properties

3.1 Identifying the State Space

It is easy to see that certain aspects of the representation cannot possibly be unique. For example, the second stage utility functions $U(\cdot, s)$ are unique only up to affine transformations that may depend on s. So, while it is easy to show that μ is unique holding everything else fixed, it will change if we make different affine transformations to the utility functions. Thus, the probabilities are not meaningful.

However, the state space, viewed as the set of possible second stage preferences, is meaningful, as will be apparent from two results. First, if we restrict attention to additive EU representations, there is an essentially unique state space. This result continues to hold if we drop the assumption of additivity across S but retain the requirement that all second stage preferences are EU. Second, if we drop the EU requirement and the additivity across S, there is still a natural reason to focus on EU state spaces: such state spaces are always minimal.

To state our results more precisely requires some definitions. The most general kind of representation we consider is the following.

Definition 2 *An* ordinal representation *of $\succ$ is a set S,*[21] *state-dependent utility functions $U : S \times \Delta(B) \to \mathbf{R}$, and an* aggregator *$u : \mathbf{R}^S \to \mathbf{R}$ such that $V : \mathcal{P}(\Delta(B)) \to \mathbf{R}$ represents preferences, where*

$$V(x) = u[(\sup_{\beta \in x} U(\beta, s))_{s \in S}]$$

and where u is strictly increasing on $U^*(X) \equiv \{(\sup_{\beta \in x} U(\beta, s))_{s \in S} \mid x \in X\}$.[22]

[21] As with our definition of additive EU representations (see footnote 17), to be precise, we need a measure space where U is measurable with respect to the space.

[22] It might appear that the representation in Theorem 1 does not satisfy this last requirement when the state space is uncountable and the measure is nonatomic. The requirement implies that if we compare two sets, x and x', such that x has a higher maximal expected utility at one state s ($\sup_{\beta \in x} U(\beta, s) > \sup_{\beta \in x'} U(\beta, s)$) and equal maximal expected utility at all other states

Definition 3 *An* ordinal EU representation *is an ordinal representation where each $U(\cdot, s)$ is an expected-utility function in the sense that*

$$U(\beta, s) = \sum_{b \in B} U(b, s)\beta(b).$$

Of course, our additive EU representation is an ordinal EU representation where the aggregator is additive.

Finally, we must define what we mean by the set of second-stage preferences in a representation. Given a representation of any of the above forms and a state in that representation, we define $\succ_s^*$ to be the preference relation over $\Delta(B)$ represented by the utility function $U(\cdot, s)$. That is, $\succ_s^*$ is defined by

$$\beta \succ_s^* \beta' \iff U(\beta, s) > U(\beta', s).$$

It should be clear that the only important aspect of the state space is the set of preferences it corresponds to. In light of this, we henceforth refer to the set $\{\succ_s^* | \ s \in S\}$ as the state space.

We say that a representation (of any of the above forms) is *nonredundant* if there is no state with complete indifference — that is, there is no state s such that $\beta \sim_s^* \beta'$ for all β and β'.[23] To understand this, note that the preference relation over menus cannot identify whether or not there is a state s with complete indifference as the presence of such a state simply adds a constant to the value of any x.

Our results are strongest and clearest when there is an additive EU representation with a finite state space. In this case, given the restriction to EU representations, the state space is uniquely identified. That is, every nonredundant ordinal or additive EU representation has the same state space. Once we move outside the class of EU representations to general ordinal representations, this uniqueness no longer holds. However, any

$(\sup_{\beta \in x} U(\beta, s') = \sup_{\beta \in x'} U(\beta, s')$ for all $s' \neq s)$, then x is better: $u[(\sup_{\beta \in x} U(\beta, s'))_{s'}, s] > u[(\sup_{\beta \in x'} U(\beta, s'))_{s'}, s]$. But changing the value of the function being integrated at one point will not change the integral if μ is absolutely continuous. To understand this requirement, note that if we have a continuous measure and if $U(\cdot, s)$ is continuous in s, then if $\sup_{\beta \in x} U(\beta, s)$ increases at one s, it must increase on a neighborhood, so this difficulty cannot arise. Put differently, if a point is included in the state space, then it must matter in the sense that an increase in the payoff at that state always "counts," either because μ has an atom at that point or because the state is part of a group of "nearby" states which behave similarly. Thus, while we did not state this property of "increasing on $U^*(x)$" as part of the definition for additive EU representations, the representation we constructed in the proof of Theorem 1 satisfies this property.

[23] To use this condition, we must rule out the trivial preference where $x \sim x'$ for all menus x and x'. It is not hard to show that the only state space for such preferences consists of one state with complete indifference over all of $\Delta(B)$. Hence, when we speak of a nonredundant representation, we are implicitly ruling out such trivial preferences.

ordinal representation whose state space differs from the unique nonredundant additive EU state space has a strictly larger number of states.

Outside the finite case, our results are a little more complex. The general version of our uniqueness result focuses on the closure of the state space, rather than the state space itself. To understand why, suppose $B = \{L, M, R\}$ and suppose we have an additive EU representation with $S = \{1, 2, \ldots\}$, a measure μ with full support on S, and state dependent utilities given by

$$U(L, m) = 1$$
$$U(M, m) = 1/m$$
$$U(R, m) = 0$$

There is an obvious sense in which the "limit" of this sequence of preferences is the preference represented by von Neumann–Morganstern utility function $\bar{U}$ with $\bar{U}(L) = 1$ and $\bar{U}(M) = \bar{U}(R) = 0$. Is this preference in the support? The answer would seem to be no: there is no s in the support of μ such that $\succ_s^*$ is this preference. On the other hand, suppose we rewrote the representation to replace S above with the induced set of von Neumann–Morganstern utility functions. That is, we could replace state 1 above with $(1, 1, 0)$, state 2 with $(1, 1/2, 0)$, etc. Since the support of a measure must be closed and since this sequence of vectors converges to $(1, 0, 0)$, we would now conclude that this preference must be in the state space.

In other words, we cannot meaningfully distinguish between an EU state space and its closure — any difference between them for an additive EU representation is only a matter of how we choose to write the measure. A similar issue arises for ordinal EU representations: we cannot identify the presence or absence of limit points. To see the point, suppose we have an additive EU representation with a state space given by the one in the example above together with the "limit state", say ∞, where

$$U(L, \infty) = 1$$
$$U(M, \infty) = U(R, \infty) = 0$$

Suppose the measure puts probability $(1/2)^{m+1}$ on state m and probability $1/2$ on state ∞. Then we could rewrite this as an ordinal EU representation without the "limit state" as follows. Let the aggregator, u, be defined by

$$u(w_1, w_2, \ldots) = \sum_{m=1}^{\infty} \left(\frac{1}{2}\right)^{m+1} w_m + \frac{1}{2} \lim_{m \to \infty} w_m.$$

(It is easy to see that the limit must exist for "relevant" w vectors — that is, for any w vector in $U^*(X)$.) Letting

$$V(x) = \sum_{s \in S} \left(\frac{1}{2}\right)^{m+1} \sup_{\beta \in x} U(\beta, s) + \frac{1}{2} \sup_{\beta \in x} U(\beta, \infty).$$

it is easy to see that for any x,

$$V(x) = u[(\sup_{\beta \in x} U(\beta, m))_{m \in \{1,2,\ldots\}}].$$

In other words, the aggregator can take the limit itself and act as if the limit state were included. Hence, again, it is impossible to identify the presence or absence of such limit points in the state space.[24]

We define the closure of a set of EU preferences by defining convergence of a sequence of EU preferences. This topology was also used by Dhillon and Mertens (1996). It is also not hard to show that it justifies the intuitive statements above about limits of preferences.

Definition 4 *Given a sequence $\{\succ_k^*\}$ of expected–utility preferences*[25] *over $\Delta(B)$, we say that $\succ^*$ is a limit of the sequence if it is a nontrivial*[26] *expected utility preference such that*

$$\beta \succ^* \beta' \text{ implies } \exists K \text{ such that } \beta \succ_k^* \beta', \ \forall k \geq K.$$

The closure of a set of expected–utility preferences adds to the set all such limit points.

Theorem 11 *If R_1 and R_2 are nonredundant ordinal EU representations of the same preference $\succ$ where $\succ$ satisfies continuity and monotonicity, then the closure of the state spaces for R_1 and R_2 are the same.*

This result is proved in the appendix but we provide a proof sketch for a simple special case to give the intuition. Consider the case where $B = \{L, M, R\}$ and where there are two nonredundant ordinal EU representations with finite state spaces. With finite spaces, the closure is just the state space itself, so our result says that these state spaces must be the same. Calling these representations 1 and 2 for brevity, suppose, contrary to our claim, that representation 2 has a second–stage preference, say $\succ^*$, which is not in the state space for representation 1. Fix any (interior) indifference curve, say I^*, for this preference and let x be the (weak) lower contour set for I^* (see Figure 1) — that is, the set of lotteries "below" (in utility) or on the indifference curve.

Suppose that for each second–stage preference, say $\succ_{s_1}$ in the state space of representation 1, we find the highest possible indifference curve subject to the constraint that the

[24]This is not to say that there is no way to bring in additional information about the agent which would enable us to make this distinction, only that preferences over menus cannot be used to identify the state space more precisely than up to closure.

[25]That is, preferences satisfying the von Neumann–Morganstern axioms.

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

indifference curve have a nonempty intersection with x. It is easy to see that there are three possibilities. First, the best possible point for the preference $\succ_{s_1}$ may lie in the set x, so that the constraint is not binding. In this case, the lower contour set associated with this preference and this constraint would be the entire simplex. Second, it can happen that the best point for the preference $\succ_{s_1}$ lies at the point labelled α in Figure 1. The dotted line from α would be a typical indifference curve giving this optimum. It should be viewed as the indifference curve with the slope closest to that of I^* among all the indifference curves from the preferences in representation 1 which have α as the optimal choice. Because representation 1 has a finite state space, there is a closest such indifference curve. Because that state space consists entirely of expected–utility preferences and does not contain the preference $\succ^*$, this closest indifference curve *cannot* coincide with I^*. If the indifference curves did coincide, the expected–utility property would say that the preference generating this steepest indifference curve is exactly $\succ^*$, contradicting our premise. The third possibility is the analogous case where the best point is β. Again, the dotted line from β should be interpretted as the closest such indifference curve to I^*. Again, this indifference curve cannot coincide with I^*.

The key implication of this is that if we take the intersection of these various lower contour sets, we get $x \cup y$ as shown in Figure 1. Because neither dashed indifference curve can coincide with I^*, the set y *must* be nonempty.

To see the significance of this, fix any preference in representation 1 and compare the level of utility achievable when the feasible set is x to the level feasible when it is $x \cup y$. It is easy to see that in any of the three cases above, adding y to the feasible set does not enable the agent to increase her payoff. Hence for every s in representation 1, we must have

$$\max_{\beta \in x} U(\beta, s) = \max_{\beta \in x \cup y} U(\beta, s).$$

Hence $x \sim x \cup y$. On the other hand, clearly, if the preferences are $\succ^*$, the agent would be strictly better off choosing from $x \cup y$ than choosing from x alone. Since the agent cannot be made worse off from the expansion of the feasible set, we see that for every state s in representation 2,

$$\max_{\beta \in x} U(\beta, s) \leq \max_{\beta \in x \cup y} U(\beta, s),$$

with a strict inequality for the state where the agent's preference is $\succ^*$. Hence representation 2 has $x \cup y \succ x$, a contradiction. ▌

Remark 2 It is worth noting that this result does not require the independence axiom. While Theorem 1 showed that an additive EU representation exists only if the preferences satisfy independence, a preference can have an ordinal EU representation without satisfying this axiom.

As we noted in the introduction, one reason the lack of identification of the state space in Kreps' framework is problematic is that it makes it quite difficult to relate the structure of the state space to intuitive properties of the underlying preferences. For example, a natural intuition is that larger state spaces correspond to a greater "concern" about unforeseen contingencies. In the Kreps framework, this not true: alternative representations of the *same* preference can have nested state spaces. Since we pin down the state space (given the EU restriction), a natural conjecture is that larger state spaces do correspond to more "concern" about unforeseen contingencies. In fact, our proof of Theorem 11 shows that this conjecture is correct. Specifically, say that $\succ_2$ *values flexibility more* than $\succ_1$ if

$$x \cup \{\beta\} \succ_1 x \text{ implies } x \cup \{\beta\} \succ_2 x.$$

In other words, whenever $\succ_1$ values the greater flexibility of $x \cup \{\beta\}$ over x, $\succ_2$ does as well, so the latter values flexibility at least as much as the former. Intuitively, in this approach, the value of flexibility is attributed to unforeseen contingencies. Hence an agent who values flexibility more than another does so because she is more concerned with unforeseen contingencies. The following theorem says that if one agent values flexibility more, then, as we'd expect, the closure of her state space must be larger.

Theorem 12 *Let R_i denote a nonredundant ordinal EU representation of preference $\succ_i$, $i = 1, 2$, where both preference relations satisfy monotonicity and continuity. Suppose $\succ_2$ values flexibility more than $\succ_1$. Then the closure of the state space for R_1 is a subset of the closure of the state space for R_2.*

It is easy to see that Theorem 11 is a corollary to this result. In short, our identification of the state space enables us to relate its structure to intuitive, economic properties of the underlying preferences.

Thus given the restriction to expected utility preferences in the second stage, we obtain the strong result that the set of second stage preferences is essentially unique. If we drop the restriction to second–stage expected–utility preferences, we still have a natural reason for focusing attention on expected–utility preferences: such state spaces are minimal. As noted above, in the finite case, our result is that any ordinal representation whose state space is not the same as the unique EU state space must have a strictly larger state space. In the infinite case, the result is more complex and requires us to define a kind of representation in between additive and ordinal representations. This representation, which we call a *finitely additive EU representation*, is identical to an additive EU representation, except we only require μ to be a finitely additive measure rather than a standard (countably additive) measure.

We view finitely additive EU representations as slightly less appealing than additive EU representations. Finitely additive measures are undeniably less convenient for characterization purposes than countably additive ones and can lead to odd properties, as discussed by Stinchcombe (1997). In our case, the kinds of paradoxes explained by Stinchcombe cannot arise,[27] however, so we do not see this representation as significantly less desirable than the simpler additive EU case.

In the infinite state case, our result is that there is always an ordinal EU representation, either additive or finitely additive, whose state space has the smallest cardinality among all state spaces of all ordinal representations. More precisely,

Theorem 13 *Given our axioms:*

1. *If there is a nonredundant ordinal EU representation with a finite state space, then any ordinal representation, EU or otherwise, with a different state space has a strictly larger one.*

2. *If there is a nonredundant ordinal EU representation with an infinite state space, then every ordinal representation has an infinite state space. Hence if there is a nonredundant ordinal EU representation with a countably infinite state space, the state space for every ordinal representation must have weakly larger cardinality.*

3. *There is always a finitely additive EU representation with a countable state space.*

Again, the proof is contained in the appendix but the basic intuition is not difficult. For simplicity, we describe this intuition only for the case where there is a nonredundant ordinal EU representation, say representation 1, with a finite state space. Fix any preference, say $\succ^*$, in the state space of representation 1 and any (interior) lower contour set, say x, for that preference. Fix any set of lotteries y which is disjoint from x. If the agent's preferences are given by $\succ^*$, then clearly she is strictly better off choosing from $x \cup y$ than from x alone. After all, x is the set of lotteries yielding utility less than some amount according to $\succ^*$, so everything in y must yield higher utility. Since the agent cannot be worse off choosing from $x \cup y$ than from x in any other state, this tells us 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? Clearly, 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

[27] The reason finite additivity causes no problems for us is that it only allows us to construct a smaller state space, not to represent otherwise unrepresentable preferences. Hence the preferences must be just as well behaved under a finitely additive EU representation as under the regular additive EU.

infinite state space, then our minimality property clearly holds, so we need not consider this case further.)

In light of this, fix any interior β and consider the collection of lower contour sets for the different preferences in the state space for representation 1. Since these are all expected utility preferences, each different state must be associated with a different lower contour set. But each of these lower contour sets must be associated with a different second stage preference in representation 2. Hence representation 2 must have at least as many possible second-stage preferences as representation 1. In fact, the proof of Theorem 13 shows that this comparison must be strict unless the state spaces are the same. ∎

In short, we see that our axioms imply that there is an additive (or finitely additive) EU representation whose state space has the minimum possible cardinality. In this sense, such representations have the "simplest" possible state space.

3.2 Meaningful Additivity

As discussed in the introduction, the additivity of Kreps' representation is not meaningful in the sense that it does not imply any restriction on the preferences. That is, in his framework, preferences have an additive representation whenever they have an ordinal one. It is true that some ordinal representations are intrinsically nonadditive in the sense that the aggregator is not a monotone transformation of a summation. Still, preferences with such ordinal representations do have an additive representation but on a different state space.[28]

As seen above, we can pin down the state space in our framework, suggesting that additivity is more meaningful here. In this subsection, we demonstrate that this is correct: additivity is an implication of our axioms. More specifically, the aggregator for any ordinal EU representation of preferences satisfying our axioms is a monotone transformation of an integral and so is intrinsically additive. Hence any preference with an ordinal EU representation which is intrinsically nonadditive is excluded by our axioms. Since our axioms are necessary and sufficient for the existence of an additive EU representation, this implies that for such preferences, we cannot simply change the state space and find an additive EU representation. In this sense, additivity of the representation is meaningful here in a way it is not in Kreps (1979).

This intuition suggests that a distinction between ordinal and additive might be obtained in Kreps' model by means of a restriction on the second-stage preferences, analogous to the way we require expected-utility in the second stage. While it is conceivable

[28]See Kreps (1979), page 573.

that there is some restriction on the set of preferences over menus such that this could be done, it cannot be done within the class of preferences Kreps studies. In Section C of the appendix, we show by example that there is *no* restriction on second-stage preferences possible in Kreps' framework which (a) allows an additive representation of every preference he considers and (b) does not allow any intrinsically nonadditive ordinal representation.

On the other hand, we do satisfy these two criteria for the set of preferences we consider. Theorem 1 showed that every preference satisfying our axioms has an additive EU representation. The following result shows that none of them has an intrinsically nonadditive ordinal EU representation.

Theorem 14 *Let (S, U, u) be an ordinal EU representation of any preference satisfying monotonicity, continuity, and independence. Then there exists a finitely additive measure μ on S such that for any $x \subseteq \Delta(B)$,*

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

up to a monotone transformation.

It is important to note that this result does *not* just say that given any ordinal EU representation, there is some (finitely) additive EU representation that is a monotone transformation of it. 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 the aggregator function u must be a monotone transformation of an integral (at least restricted to the relevant set of points in $\mathbf{R}^S$). In this sense, all ordinal EU representations are intrinsically additive.

The intuition of the proof, contained in the Appendix, is very simple. From Proposition 3, we know that, up to a monotone transformation, we must be able to write

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

for some function V which is linear in the sense that

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

Also, it is not hard to use the fact that $U(\cdot, s)$ is an expected–utility function to show that for all $s \in S$,

$$\sup_{\beta \in \lambda x + (1-\lambda)x'} U(\beta, s) = \lambda \sup_{\beta \in x} U(\beta, s) + (1 - \lambda) \sup_{\beta \in x'} U(\beta, s).$$

Hence up to a monotone transformation,

$$u((\lambda \sup_{\beta\in x} U(\beta,s) + (1-\lambda)\sup_{\beta\in x'} U(\beta,s))_{s\in S})$$
$$= \lambda u((\sup_{\beta\in x} U(\beta,s))_{s\in S}) + (1-\lambda) u((\sup_{\beta\in x'} U(\beta,s))_{s\in S}).$$

Therefore, u must be (extendable to) a linear functional and hence can be written as an integral with respect to a finitely additive measure.

4 Conclusion

To summarize, we have extended Kreps (1979) in four important ways. First, we have shown how additional structure can be put on the second–stage preferences. Because the second–stage preferences are only indirectly identified through preferences for flexibility, this is not a standard exercise. Second, using this additional structure, we showed that the state space is (essentially) unique within the class of representations we primarily focus on. That is, there is an essentially unique state space which must be used in every nonredundant EU representation, additive or not. Third, even in a much broader class of representations, our additional structure implies that the EU state space is minimal and thus a natural focus of study. Finally, because we are able to restrict the state space this way, the additivity of the representation is meaningful. As illustrated in Theorem 12, pinning down the 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.

A A review of the Hausdorff topology

Let d denote any distance on $\Delta(B)$. For any pair $x, x' \in \mathcal{P}(\Delta(B))$, we define as usual

$$d(\alpha, x') \equiv \inf_{\beta \in x'} d(\alpha, \beta); \tag{5}$$

and

$$e(x, x') \equiv \sup_{\alpha \in x} d(\alpha, x'), \tag{6}$$

the *excess* of x over x'. As is well known, one can define two hemimetrics over $\mathcal{P}(\Delta(B))$ with this function. In particular the ball around x of radius ϵ in the upper hemimetric topology is the set:

$$\mathcal{B}^u(x, \epsilon) \equiv \{x' \in \mathcal{P}(\Delta(B)) \mid e(x, x') < \epsilon\} \tag{7}$$

while the ball in the lower hemimetric topology is the set:

$$\mathcal{B}^l(x, \epsilon) \equiv \{x' \in \mathcal{P}(\Delta(B)) \mid e(x', x) < \epsilon\}; \tag{8}$$

and finally the ball in the hemimetric topology is the set:

$$\mathcal{B}(x, \epsilon) \equiv \{x' \in \mathcal{P}(\Delta(B)) \mid \max\{e(x', x), e(x, x')\} < \epsilon.\} \tag{9}$$

Clearly $\mathcal{B}(x, \epsilon) = \mathcal{B}^u(x, \epsilon) \cap \mathcal{B}^l(x, \epsilon)$. The topology whose basis is the balls as defined in (9) is the Hausdorff hemimetric topology.

B Proofs

Proof of Lemma 2. By monotonicity, $\mathrm{cl}(x) \succeq x$. So we only have to show that $\mathrm{cl}(x) \succ x$ is impossible. To see this, suppose that it does hold. By continuity, then, for some $\epsilon > 0$

$$\mathrm{cl}(x) \succ x', \text{ for every } x' \in \mathcal{B}(x, \epsilon). \tag{10}$$

Now define:

$$x^\epsilon \equiv \{\beta \in \Delta(B) \mid d(\beta, x) < \epsilon\} = \{\beta \in \Delta(B) \mid d(\beta, \beta') < \epsilon \text{ for some } \beta' \in x\}.$$

Clearly, $\mathrm{cl}(x) \subseteq x^\epsilon$, so $x^\epsilon \succeq \mathrm{cl}(x)$ by monotonicity.

But by definition of x^ϵ, $\sup_{\beta \in x^\epsilon} d(\beta, x) < \epsilon$. Also, though, x^ϵ is the set of β such that $d(\beta, \beta') < \epsilon$ for some $\beta' \in x$, so $\sup_{\beta' \in x} d(\beta, x) < \epsilon$. Hence $x^\epsilon \in \mathcal{B}(x, \epsilon)$, implying $\mathrm{cl}(x) \succ x^\epsilon$, a contradiction.

We now show that $\text{conv}(x) \sim x$. First we show the result for any finite x. This step only uses monotonicity and independence.

So suppose $\text{conv}(x) \not\sim x$. By monotonicity, we must have $\text{conv}(x) \succ x$. By independence, for every $\lambda \in (0, 1)$,

$$\lambda \text{conv}(x) + (1-\lambda)\text{conv}(x) \succ \lambda x + (1-\lambda)\text{conv}(x).$$

It's easy to see that the set on the left–hand side must be $\text{conv}(x)$. To see this, note that $x' \subseteq \lambda x' + (1-\lambda)x'$ for any set x' and any $\lambda \in [0, 1]$. For the converse, simply observe that if $\beta = \lambda\beta_1 + (1-\lambda)\beta_2$ where $\beta_1, \beta_2 \in \text{conv}(x)$, then β must also be a convex combination of elements of x and hence must also be in $\text{conv}(x)$. Hence we have

$$\text{conv}(x) \succ \lambda x + (1-\lambda)\text{conv}(x)$$

for all $\lambda \in (0, 1)$. Let $x = \{\beta_1, \ldots, \beta_n\}$ (recall that x is finite). We now show that for all $\lambda \in (0, 1/n]$, $\lambda x + (1-\lambda)\text{conv}(x) = \text{conv}(x)$, yielding a contradiction.

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 $\lambda \in (0, 1/n]$ and any $\beta \in \text{conv}(x)$. By definition, there are nonnegative numbers t_i, $i = 1, \ldots, n$, such that $\sum_i t_i = 1$ and $\sum_i t_i\beta_i = \beta$. Clearly, there must be some j such that $t_j \geq 1/n$. Define $\hat{t}_i$ for $i = 1, \ldots, n$ 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, $t_j \geq 1/n \geq \lambda$ 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 $\hat{\beta} = \sum_i \hat{t}_i\beta_i$. 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 $\lambda\beta_j + (1-\lambda)\hat{\beta} = \sum_i t'_i\beta_i$ 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 $\lambda x + (1-\lambda)\text{conv}(x) = \text{conv}(x)$.

So we get $\text{conv}(x) \succ \text{conv}(x)$, a contradiction. Hence for every finite x, $x \sim \text{conv}(x)$. So fix any set x such that $\text{conv}(x)$ has finitely many extreme points. Let x^* be the set of extreme points. Then we must have

$$x^* \subseteq x \subseteq \text{conv}(x),$$

so monotonicity implies $x^* \preceq x \preceq \text{conv}(x)$. By the above, though, $x^* \sim \text{conv}(x)$, so $x \sim \text{conv}(x)$. This completes the proof for finite x.

Next we use continuity to extend this to any closed x which is sufficient for the result by the fact established above that $x \sim \text{cl}(x)$. First, we prove that for every closed $x \subseteq \Delta(B)$, there is a countable set $E(x)$ such that $\text{conv}(E(x)) = \text{conv}(x)$.

It is well known that the support function for a closed x, $\sigma : \mathbf{R}^B \to \mathbf{R}$, is continuous and satisfies

$$\text{conv}(x) = \bigcap_{p \in \mathbf{R}^B} H(p), \tag{11}$$

where

$$H(p) \equiv \{y \in \mathbf{R}^B \mid (p, y) \leq \sigma(p)\},$$

where (p, y) is the inner product. Hence, letting D be the set of elements in $\mathbf{R}^B$ with rational coordinates,

$$\text{conv}(x) = \bigcap_{p \in D} H(p). \tag{12}$$

To see this, note that $D \subset \mathbf{R}^B$ implies $\cap_{p \in \mathbf{R}^B} H(p) \subseteq \cap_{p \in D} H(p)$. To see that the inclusion cannot be strict, consider any $\beta \notin \cap_{p \in \mathbf{R}^B} H(p)$. Then there is some $p \in \mathbf{R}^B$ such that $(p, \beta) > \sigma(p)$. By continuity of σ and density of D, the same inequality obtains for a $p' \in D$, implying $\beta \notin \cap_{p \in D} H(p)$. Hence $\cap_{p \in \mathbf{R}^B} H(p) = \cap_{p \in D} H(p) = \text{conv}(x)$.

For each $p \in D$, choose any $\xi(p) \in \text{argmax}_{\beta \in x}(p, \beta)$. (This is well defined because x is compact.) Let $E(x) = \{\xi(p) \mid p \in D\}$. Obviously, $E(x)$ is countable. Then

$$\text{conv}(x) = \text{conv}(E(x)).$$

To see this, note that the inclusion $\text{conv}(\{\xi(p)\}_{p \in D}) \subseteq \text{conv}(x)$ is immediate from the definition of ξ; that this inclusion cannot be strict is easily shown by arguing by contradiction and using the separation theorem.

To conclude the proof we combine these two results, namely that for any finite x, $\text{conv}(x) \sim x$ for finite x, and that for any closed x, $\text{conv}(x) = \text{conv}(E(x))$ for a countable set $E(x)$. So for $x \subseteq \Delta(B)$, take a sequence e^n of finite subsets of $E(x)$, such that $e^n \subset e^{n+1}$ for every n, and $\cup_{n \geq 1}(e^n) = E(x)$; then define $x_n \equiv \text{conv}(e_n)$. Clearly $x^n \subseteq x$ for every n, so:

$$x \succeq x^n \sim \text{conv}(x^n),$$

from monotonicity and the result for finite x. Also, $\text{conv}(x^n) \to \text{conv}(x)$, so $x \succeq \text{conv}(x)$ from the continuity axiom. ▌

Proof of Proposition 3: We make use of the following lemma.

Lemma 15 *If preferences satisfy continuity and monotonicity, then there is a continuous $V_D : \hat{X} \to \mathbf{R}$ that represents preferences, i.e., such that $x \succ x'$ iff $V(x) > V(x')$.*

Proof. Note that $\Delta(B)$ is connected, compact, and metric. Hence (X, τ_H), where τ_H is the Hausdorff topology, is:

1. separable (see Theorem 4.5.5, page 51, of Klein and Thompson (1984));

2. connected (see Theorem 2.4.6, page 20, of Klein and Thompson (1984)).

Now all the conditions of Debreu's theorem (see for instance Fishburn (1970), Lemma 5.1, page 62) are satisfied, giving the desired representation. ▌

Next, we verify that the mixture space axioms (see Kreps (1988), page 52) hold for X. The only mixture space condition which is not trivial to verify is the Herstein Milnor continuity condition. We now show that our continuity condition implies that if $x, x', x'' \in \hat{X}$ 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 $\lambda, \mu \in [0, 1]$,

$$\begin{aligned} d_{\text{Hausdorff}}(x(\lambda), x(\mu)) &= d_{\text{supnorm}}(\sigma_{x(\lambda)} - \sigma_{x(\mu)}) \\ &= \mid \lambda - \mu \mid d_{\text{supnorm}}(\sigma_x - \sigma_{x''}) \end{aligned}$$

where the first equality is part (4) of Lemma 4, and the second follows from part (1) of Lemma 4. Hence the function from $[0, 1]$ to $\hat{X}$ with the Hausdorff topology defined by $x(\lambda)$ is continuous. Now the result follows from the continuity of V_D (see Lemma 15).

Remark 3 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.

Hence by the Herstein Milnor theorem, there is a linear V which represents the preferences and is unique up to an affine transformation. By this uniqueness and Lemma 15, V must be continuous in the Hausdorff topology. ▌

Proof of Lemma 7: 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 $C(S^n)$;

2. H^* separates the points of S^n;

3. H^* contains the constant function 1_{S^n}.

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 $\cup_{r>0}C_+$ and C_+ is convex and contains the zero function.[29] Lemma 5 implies that H contains the supremum of any two of its elements. Next, note that $C(S^n)$ 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 $C(S^n)$).[30]

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 infimun of two of its elements. While H is not closed under taking infimums, this follows for $H^* = H - H$ by taking negatives. Specifically, $(\sigma_1 - \sigma_2) \wedge (\sigma_1' - \sigma_2') = -[(\sigma_2 - \sigma_1) \vee (\sigma_2' - \sigma_1')] = -[((\sigma_1 + \sigma_2') \vee (\sigma_1' + \sigma_2)) - (\sigma_2 + \sigma_2')] = (\sigma_2 + \sigma_2') - [(\sigma_1 + \sigma_2') \vee (\sigma_1' + \sigma_2)]$. Now repeat the preceding argument. ▌

Step (2). Let $s, s' \in S^n, s \neq s'$. Note first that for any $x \in \hat{X}$ which contains $(1/n, \ldots, 1/n)$, 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, $\lambda\alpha + (1-\lambda)(1/n, \ldots, 1/n) \equiv \alpha(\lambda) \in \Delta(B)$, and if we let $x \equiv \{\theta\alpha(\lambda) + (1-\theta)(1/n, \ldots, 1/n) \mid \theta \in [0, 1]\}$ then we have: $\sigma_x(s) = \lambda(s, \alpha) > \sigma_x(s')$ as claimed. ▌

Step (3). Follows from Lemma 5 and the definition of H. ▌

[29] 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 $f_i = r_i g_i, g_i \in C_+, i = 1, 2$; say $r_2 \leq r_1$. So if $\lambda \in [0, 1]$ then $\lambda f_1 + (1-\lambda) f_2 = \lambda r_1 g_1 + (1-\lambda)\frac{r_2}{r_1} r_1 g_2 = r_1[\lambda g_1 + (1-\lambda)\frac{r_2}{r_1} g_2]$. But the function in square brackets is in C_+ because $\frac{r_2}{r_1} g_2 \in C_+$. This last statement follows in turn because $\frac{r_2}{r_1} \in (0, 1)$, and C_+ is a convex set which contains the zero.

[30] We defined the supremum, $\sigma \vee \sigma'$, in part (4) of Lemma 5; the infinimum, denoted $\sigma \wedge \sigma'$, is defined similarly. For $\sigma, \sigma' \in C(S^n)$, and for $r \in \mathbf{R}$, addition, $\sigma + \sigma'$, and scalar multiplication, $r\sigma$, are both defined in the usual way, under which $C(S^n)$ 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.

Proof of Lemma 8: We have seen that H^* is a subspace of $C(S^n)$, and that W is a real linear functional on H^*. For any $f \in H^*$, $- \parallel f \parallel 1_{S^n} \leq f \leq \parallel f \parallel 1_{S^n}$, where $\parallel f \parallel$ is the supremum norm in $C(S^n)$ and 1_{S^n} is the indicator function of S^n. Since W is monotonic on H^*, $-W(1_{S^n}) \parallel f \parallel \leq W(f) \leq W(1_{S^n}) \parallel f \parallel$, that is $|W(f)| \leq W(1_{S^n}) \parallel f \parallel$. Hence by the Hahn Banach theorem (see Theorem 4, page 187 of Royden (1968)), W has an extension to a continuous linear functional on $C(S^n)$. That this extension is unique follows from the fact that H^* is dense in the supremum norm in $C(S^n)$, as shown in Lemma 7. ▌

Proof of Lemma 10: By Lemma 2 and Proposition 9, for any $x \in \mathcal{P}(\Delta(B))$,

$$V(x) = \int_S \max_{\beta \in \text{conv}(\text{cl}(x))} U(\beta, s)\, \mu(ds).$$

Fix any closed $x \in \mathcal{P}(\Delta(B))$ and any $s \in S$. Clearly, $x \subseteq \text{conv}(x)$ implies

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

The linearity of $U(\cdot, s)$ implies that this inequality can never be strict. Hence for every $x \in \mathcal{P}(\Delta(B))$,

$$V(x) = \int_S \max_{\beta \in \text{cl}(x)} U(\beta, s)\, \mu(ds) = \int_S \sup_{\beta \in x} U(\beta, s)\, \mu(ds).$$

▌

Proofs of Theorems 11 and 12: These proofs and the proof of Theorem 13 make use of a proposition which 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 $\mathbf{P}_R$ denote the set of preferences in the representation that is,

$$\mathbf{P}_R = \{\succ_s^* \mid s \in S\}.$$

For any preference $\succ^*$ over Δ and any $\beta \in \Delta$, let

$$\mathcal{L}_{\succ^*}(\beta) = \{\beta' \in \Delta \mid \beta \succeq^* \beta'\}.$$

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

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

Lemma 16 *For any ordinal representation R of $\succ$ and any $x \in \mathcal{L}_R$,*

1. $x \prec x \cup \{\beta\}$ *for any* $\beta \notin x$.

2. x *is closed if* $\succ$ *satisfies monotonicity and continuity and is convex if* $\succ$ *also satisfies independence.*

Proof of Lemma.

Part (*i*): Suppose $x = \mathcal{L}_{\succ_s^*}(\beta')$ for $\succ_s^* \in \mathbf{P}_R$. By the definition of a lower contour set, for all $\beta \notin x$,

$$\max_{\hat{\beta} \in x} U(\hat{\beta}, s) = U(\beta', s) < U(\beta, s) = \max_{\hat{\beta} \in x \cup \{\beta\}} U(\hat{\beta}, s).$$

Since enlarging a set cannot make the agent worse off in any state, the fact that the aggregator is strictly increasing on the relevant set implies that $x \cup \{\beta\} \succ x$.

Part (*ii*): Suppose $x \in \mathcal{L}_R$ is not closed. Let β be a point in its closure such that $\beta \notin x$. By part (i), we must have $x \cup \{\beta\} \succ x$, contradicting Lemma 2. Suppose $x \in \mathcal{L}_R$ is not convex. Then there exists $\beta = \lambda\beta_1 + (1-\lambda)\beta_2$, where β_1 and β_2 are in x, but β is not. Since $\beta \notin x$, (i) implies $x \cup \{\beta\} \succ x$. But $x \sim \text{conv}(x) = \text{conv}(x \cup \{\beta\}) \sim x \cup \{\beta\}$ by Lemma 2, a contradiction. ▌

Proposition 17 *Let* $R_i = (S_i, U_i, u_i)$, $i = 1, 2$, *denote ordinal representations of preferences* $\succ_i$, *where both preference relations satisfy monotonicity and continuity. Suppose* $\succ_2$ *values flexibility more than* $\succ_1$ *in the sense that*

$$x \cup \{\beta\} \succ_1 x \quad \text{implies} \quad x \cup \{\beta\} \succ_2 x.$$

Then for every 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* $x \in \mathcal{L}_{\succ^*}$*, there is an index set* K*, a nonrepeating sequence of states in representation* j*,* $\{s_k\}_{k \in K} \subseteq S_2$*, and a sequence of lower contour sets* $\{x_k\}_{k \in K}$ *such that* $x_k \in \mathcal{L}_{\succ_{s_k}^*}$ *and*

$$x = \bigcap_{k \in K} x_k.$$

Proof of Proposition. Let x denote any element of $\mathcal{L}_{R_1}$. For each $s \in S_2$, let

$$x_s = \{\beta' \in \Delta(B) \mid U_2(\beta', s) \leq \max_{\beta \in x} U_2(\beta, s)\}.$$

By part (ii) of Lemma 16, x must be closed so this is well defined. Clearly, each $x_s \in \mathcal{L}_{R_2}$. 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, $x \subseteq x \cup \{\beta'\} \subseteq x'$, so by monotonicity, $V_2(x) \leq V_2(x \cup \{\beta'\}) \leq V_2(x')$. 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 we must have $V_2(x') \leq V_2(x)$. Hence $x \sim_2 x \cup \{\beta'\}$. However, by Lemma 16, part (i), the fact that $x \in \mathcal{L}_{R_1}$ implies $x \cup \{\beta'\} \succ_1 x$, so the assumption that $\succ_2$ values flexibility more than $\succ_1$ implies $x \cup \{\beta'\} \succ_2 x$, a contradiction. Hence $x = x'$. ▌

We next prove Theorem 12. Theorem 11 is a corollary.

So fix any $s_1 \in S_1$. We show by contradiction that $\succ^*_{s_1} \in \text{cl}(\{\succ^*_s | \ s \in S_2\})$. So suppose not. Fix any interior β and let $x = \mathcal{L}_{\succ^*_{s_1}}(\beta)$. Note that the lower contour sets in an ordinal EU representation are half spaces intersected with $\Delta(B)$; this follows immediately from the fact that indifference curves with EU preferences are straight lines. We therefore call them half spaces. So x is a half space. By Proposition 17, x must be the intersection of some collection of lower contour sets in $\mathcal{L}_2 \equiv \mathcal{L}_{R_2}$. By supposition, $\succ^*_{s_1}$ is not in the closure of representation 2's state space. Since all the preferences are expected–utility, this means that $x \notin \mathcal{L}_2$. Since all the lower contour sets in $\mathcal{L}_2$ are half spaces also, x must be the intersection of an infinite sequence of distinct lower contour sets in $\mathcal{L}_2$. That is, because x is a half space, no finite intersection of convex sets (half spaces or otherwise) not including x could equal x. Hence there is a nonrepeating sequence of states in S_2, $\{s^2_k\}$, and a nonrepeating sequence $\{x_k\}$ with $x_k \in \mathcal{L}_{\succ^*_{s^2_k}}$ such that

$$x = \bigcap_{k=1}^{\infty} x_k.$$

Without loss of generality, we can choose this sequence to be of nested sets, so for every N, $\cap_{k=1}^{N} x_k$ is itself an element of $\mathcal{L}_2$. Hence x must be the limit of the sequence $\{x_k\}$. Because all these preferences are expected–utility preferences, it is easy to see that every lower contour set for the $\succ^*_{s_1}$ preferences can be generated as a sequence of lower contour sets from the same sequence of states $\{s^2_k\}$. This implies that $\succ^*_{s_1}$ is the limit of $\{\succ^*_{s^2_k}\}$. To see this, simply note that $\succ^*_{s_1}$ is not trivial by nonredundancy. Furthermore, if $\beta \succ^*_{s_1} \beta'$, then there is a lower contour set, say x, for $\succ^*_{s_1}$ which includes β' but not β. For k sufficiently large, the lower contour set for $\succ^*_{s^2_k}$ which converges to x must also include β' and not β. Hence for k sufficiently large, $\beta \succ^*_{s^2_k} \beta'$. By definition, then, $\succ^*_{s_1}$ is the limit of $\{\succ^*_{s^2_k}\}$. Hence $\mathbf{P}_{R_1} \subseteq \text{cl}(\mathbf{P}_{R_2})$, so $\text{cl}(\mathbf{P}_{R_1}) \subseteq \text{cl}(\mathbf{P}_{R_2})$. ▌

Proof of Theorem 13.

Part (i): Fix a nonredundant ordinal EU representation and let $\mathbf{P}$ denote its state space where $\mathbf{P}$ is finite. Let (S^o, U^o, u^o) be any nonredundant ordinal representation and let $\mathbf{P}^o$ be its state space. Clearly, the result holds if $\mathbf{P}^o$ is infinite, so assume it is also finite.

We first construct a function $f : \mathbf{P} \times \text{int}(\Delta) \to \mathbf{P}^o$ (where $\text{int}(\Delta)$ is the interior of Δ). To construct it, fix any $\succ^* \in \mathbf{P}$ and any $\beta \in \text{int}(\Delta)$. Because $\mathbf{P}$ is a nonredundant collection of EU preferences and because β is interior, $\mathcal{L}_{\succ_1^*}(\beta) \not\subseteq \mathcal{L}_{\succ_2^*}(\beta)$ whenever $\succ_1^* \neq \succ_2^*$.

For any $\succ^* \in \mathbf{P}$, the same argument as in the proof of Theorem 11 shows that $\mathcal{L}_{\succ^*}(\beta)$ must either be a lower contour set for some preference $\succ_o^* \in \mathbf{P}^o$ or else be the intersection of a family of lower contour sets for preferences in $\mathbf{P}^o$. Suppose the former case does not hold. Then there exists an index set K, a nonrepeating sequence of states $s_k \in S^o$ for $k \in K$, and lower contour sets x_k for preference $\succ_{s_k}^*$ for $k \in K$ such that

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

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 K must be smaller than the cardinality of $\mathbf{P}^o$. Recall from part (ii) of Lemma 16 that each x_k must be convex. Hence, just as in the proof of Theorem 11, the fact that $x_k \neq \mathcal{L}_{\succ^*}(\beta)$ for all k implies that K must be infinite. But this contradicts the assumption that $\mathbf{P}^o$ is finite. Hence it must be true that there is some $\succ_o^* \in \mathbf{P}^o$ such that $\mathcal{L}_{\succ^*}(\beta) \in \mathcal{L}_{\succ_o^*}$. Let $f(\succ^*, \beta)$ be any such $\succ_o^*$.

We claim that for every β, $f(\cdot, \beta)$ is one–to–one. To see this, recall that for any $\beta \in \text{int}(\Delta)$, none of the $\mathcal{L}_{\succ^*}(\beta)$ sets is contained in any other. Hence there could not be a preference relation, expected-utility or otherwise, which has more than one of these sets as a lower contour set. Also, by construction, for every $\succ^* \in \mathbf{P}$, $\mathcal{L}_{\succ^*}(\beta)$ 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 that 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 $f(\succ^*, \beta) = \bar{f}(\succ^*)$ for every β, then $\succ^*$ has the same lower contour sets as $\bar{f}(\succ^*)$, implying $\succ^* = \bar{f}(\succ^*)$. 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 Δ 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 $\mathcal{L}_{\succ^*}(\beta') \in \mathcal{L}_{\succ_o^*}$. We now show that this cannot occur.

To see this, note first that the ordinal EU representation has parallel indifference curves so the indifference curve for 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 along this line, the indifference curve through β for preference $\succ_\circ^*$ must be the indifference curve through β for some preference $\succ^* \in \mathbf{P}$. Since $\mathbf{P}$ is finite, this means that 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' which is a distance less than ε_β away from β, the slope of the $\succ_\circ^*$ indifference curve through β is the same as the slope of the $\succ_\circ^*$ indifference curve through β'. Otherwise, we would have two points arbitrarily close together with indifference curves through them whose slopes are bounded away from one another. Hence the indifference curves would intersect, a contradiction. Let d be the infimum of the distance from β^1 along this line to a point β' where the slope of the $\succ_\circ^*$ indifference curve through β' differs from the slope at β^1. Since β^2 is such a point, we know 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 $\succ_\circ^*$ 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 must equal the slope through β^2, a contradiction.

Part (ii): It is easy to see that the first part of the proof of Part (i) above shows that if there is an ordinal representation with a finite state space, then the state space for any ordinal EU representation must be of smaller cardinality. Hence if there is no ordinal EU representation with a finite state space, there cannot be any ordinal representation with a finite state space.

Part (iii): In Theorem 1, we showed that the order $\succ$ has a representation where x is evaluated by $\int_{S^n} \sigma_x(s)\mu(ds)$ for a function σ_x which is continuous in s. Fix any countable dense subset $\{s_i\}_{i=1}^\infty$ of S^n, and define the operator T from $C(S^n)$ to ℓ^∞ (the Banach space of the bounded real valued sequences, with the supremum norm), as

$$T(f) \equiv (f(s_i)_{i=1}^\infty). \tag{13}$$

It is easy to see that T is linear, continuous, of norm 1, and injective. Hence T^{-1} also exists, and is linear, continuous, and of norm 1.

We now define a correspondence T_* from regular measures on S^n to $(\ell^\infty)^*$, the dual of ℓ^∞, in two steps. First, given the measure μ on S^n, we define for every $x \in T(C(S^n))$:

$$(x, T_*\mu) \equiv \int_{S^n} (T^{-1}(x))(s)\mu(ds)$$

where $(x, T_*\mu)$ denotes the inner product. To see that $T_*\mu$ is uniquely defined, note that T^{-1} is well defined because T is injective. Also, it is easy to show that T^{-1} is a linear functional, making the right hand side a continuous linear functional on $T(C(S^n))$. Clearly,

$$|(x, T_*\mu)| \leq \| x \|_{\ell^\infty} \tag{14}$$

for any $x \in T(C(S^n))$, which is a linear subspace of ℓ^∞. Hence by the Hahn Banach theorem, there is an extension of $T_*\mu$ to a continuous linear functional on ℓ^∞, which we denote μ_*. Then we see that the order $\succ$ has a representation where x is evaluated by the function $(T(\sigma_x), \mu_*)$. Because the linear functionals on ℓ^∞ are (up to an isometric isomorphism) integration with respect to a finitely additive measure (see, for example, Dunford and Schwartz (1958), Theorem IV.8.16), μ_* is a finitely additive measure. ▌

Proof of Theorem 14. Fix an ordinal EU representation (S, U, u) of preferences satisfying monotonicity, continuity, and independence. For each x, define a function $U^*(x) : S \to \mathbf{R}$ by $U^*(x)(s) = \sup_{\beta \in x} U(\beta, s)$ and let $U^*(X)$ be the set of $U^*(x)$ functions. Letting V be the function from Proposition 3 (where we continue to use the normalization that $V(\{1/n, \ldots, 1/n\}) = 0$), we must have, up to a monotone transformation, $V(x) = u(U^*(x))$ for all x. Since the theorem allows for such monotone transformations, we assume u satisfies this equality. Also, since u is only relevant on $U^*(X)$, we view it as a function from $U^*(X)$ to the reals and will extend it (indirectly) to a larger space below.

Without loss of generality, we assume that $\min_{b \in B} U(b, s) \neq \max_{b \in B} U(b, s)$ for all s.[31] Define functions $\gamma : S \to \mathbf{R}_+$ and $\eta : S \to \mathbf{R}$ by

$$\gamma(s) \max_{b \in B} U(b, s) + \eta(s) = 1$$

and

$$\gamma(s) \frac{1}{n} \sum_{b \in B} U(b, s) + \eta(s) = 0.$$

It is easy to see that γ and η are well defined and that $\gamma(s) \neq 0$ for all s. Let

$$\bar{U}(\beta, s) = \gamma(s) U(\beta, s) + \eta(s).$$

Similarly, for each x, define a function $\bar{U}^*(x) : S \to \mathbf{R}$ by

$$\bar{U}^*(x)(s) = \sup_{\beta \in x} \bar{U}(\beta, s)$$

and let $\bar{U}^*(X)$ denote the set of such functions. Finally, define $\bar{u} : \bar{U}^*(X) \to \mathbf{R}$ by

$$\bar{u}(\bar{U}) = u\left(\frac{\bar{U} - \eta}{\gamma}\right).$$

[31] If this is violated for some s, such states can be trivially "added back in" at the end of the argument.

It is easy to show that $\bar{u}$ is linear on $\bar{U}^*(X)$ in the sense that for all $\bar{U}$, $\bar{U}'$, and λ such that $\bar{U}$, $\bar{U}'$, and $\bar{U} + \lambda\bar{U}'$ are all elements of $\bar{U}^*(X)$,

$$\bar{u}(\bar{U} + \lambda\bar{U}') = \bar{u}(\bar{U}) + \lambda\bar{u}(\bar{U}').$$

To see this, first note that for all x, x', and $\lambda \in [0, 1]$,

$$u(U^*(\lambda x + (1-\lambda)x') = \lambda u(U^*(x)) + (1-\lambda)u(U^*(x'))$$

from the fact that $V(x)$ satisfies this and that $V(x) = u(U^*(x))$. But the fact that $U(\beta, s)$ is an expected utility function implies

$$\begin{aligned} U^*(\lambda x + (1-\lambda)x')(s) &= \sup_{\beta \in \lambda x + (1-\lambda)x'} U(\beta, s) \\ &= \lambda \sup_{\beta \in x} U(\beta, s) + (1-\lambda)\sup_{\beta \in x'} U(\beta, s) \\ &= \lambda U^*(x)(s) + (1-\lambda)U^*(x')(s). \end{aligned}$$

Hence for all $U, U' \in U^*(X)$ and $\lambda \in [0, 1]$,

$$u(\lambda U + (1-\lambda)U') = \lambda u(U) + (1-\lambda)u(U').$$

It is then easy to see that the same is true for $\bar{u}$ and $\bar{U}$. Note also that, by construction,

$$\bar{U}^*(\{1/n, \ldots, 1/n\}) = 0$$

(the zero function), so the fact that $V(\{1/n, \ldots, 1/n\}) = 0$ implies $\bar{u}(0) = 0$. From here, the proof that $\bar{u}$ is linear is exactly the same as the proof that W is linear in Lemma 6, part 1.

Let $M(S)$ denote the vector space of the set of bounded measurable functions $f : S \to \mathbf{R}$. Equipping $M(S)$ with the supremum norm, denoted $\| \cdot \|$, makes it a Banach space. Also define L to be the vector subspace of finite linear combinations of functions in the set $\bar{U}^*(X)$. It is easy to see that L is a subspace of $M(S)$.[32] Extend $\bar{u}$ to the space L linearly (just as W was extended to H^* in the proof of Theorem 1 as there, the linearity of $\bar{u}$ ensures that this can be done consistently). Clearly, $\bar{u}$ is a linear functional on L.

Now we claim that for some constant C,

$$\bar{u}(\bar{U}^*(x)) \leq C \| \bar{U}^*(x) \| . \tag{15}$$

To see this, note that the fact that u and hence $\bar{u}$ must be increasing implies

$$\bar{u}(\bar{U}^*(x)) \leq \bar{u}(\| \bar{U}^*(x) \| 1_S).$$

[32] The measurability follows from our assumption that U is measurable, as explained in footnote 21.

where 1_S is a function on S which is identically one.[33] By linearity of $\bar{u}$, then

$$\bar{u}(\bar{U}^*(x)) \leq \bar{u}(\iota) \parallel \bar{U}^*(x) \parallel,$$

implying (15) with $C \equiv u(\iota)$. Given this, the linear functional $\bar{u}$ on L can be extended to a linear functional on $M(S)$ by the Hahn–Banach theorem.

Because the linear functionals on a Banach space are (up to an isometric isomorphism) integration with respect to a finitely additive measure (see, for example, Dunford and Schwartz (1958), Theorem IV.8.16), this implies that there is a finitely additive measure $\bar{\mu}$ such that for all x,

$$\bar{u}(\bar{U}^*(x)) = \int_S \bar{U}^*(x)(s)\bar{\mu}(ds).$$

Hence

$$u\left(\frac{\bar{U}^*(x) - \eta}{\gamma}\right) = \int_S \bar{U}^*(x)(s)\bar{\mu}(ds).$$

Substituting $U^*(x)$ for $\gamma\bar{U}^*(x) + \eta$ yields

$$u(U^*(x)) = \int_S [\gamma(s)U^*(x)(s) + \eta(s)]\,\bar{\mu}(ds)$$

or

$$u(U^*(x)) = \int_S U^*(x)(s)\gamma(s)\bar{\mu}(ds) + \int_S \eta(s)\bar{\mu}(ds).$$

Hence letting $\mu = \gamma\bar{\mu}$, we see that up to a monotone transformation,

$$u(U^*(x)) = \int_S \sup_{\beta \in x} U(\beta, s)\mu(ds),$$

as was to be shown. ▌

C Example

In this section, we show that restricting the set of second stage preferences in Kreps' original framework cannot generate meaningful additivity without some restriction on the set of preferences over menus he considers. We show this by example. Let $B = \{a, b, c\}$. Since preferences in Kreps are over subsets of B, not of $\Delta(B)$, the set of menus consists of the seven nonempty subsets of B. Suppose the preferences over menus are

$$\{a, b, c\} \succ \{a, b\} \sim \{a, c\} \sim \{b, c\} \succ \{b\} \sim \{c\} \succ \{a\}.$$

[33] Note that $U^*(B) = 1_S$, so the function $\parallel U^*(x) \parallel 1_S$ is an element of L. Hence u is defined at this point.

This preference has an ordinal representation with state space $S = \{s_1, s_2, s_3\}$ with utility functions given by

	s_1	s_2	s_3
a	1	0	0
b	0	1	0
c	0	0	1

To see that there must be an aggregator making this an ordinal representation, note that we only need to find an increasing function u such that

$$u(1,1,1) > u(1,1,0) = u(1,0,1) = u(0,1,1) > u(0,1,0) = u(0,0,1) > u(1,0,0).$$

and such functions clearly exist. However, any such aggregator must be intrinsicially nonadditive. To see the point, suppose there is a function $u : \{0,1\}^3 \to \mathbf{R}$ which satisfies the above and which can be written as $u(v_1, v_2, v_3) = u_1(v_1) + u_2(v_2) + u_3(v_3)$. Then $u(0,1,0) > u(1,0,0)$ implies $u_1(0) + u_2(1) > u_1(1) + u_2(0)$. Adding $u_3(1)$ to both sides, we then have $u(0,1,1) > u(1,0,1)$, which contradicts $\{b, c\} \sim \{a, c\}$. Hence this representation is intrinsically nonadditive.[34]

We claim that there is no restriction on the class of second stage preferences in Kreps that could avoid this problem. To avoid the example above, such a restriction would have to exclude at least one of the above preferences. Since the argument is completely symmetric, suppose the excluded preference is the one the agent has in state s_1, where a is strictly best and he is indifferent between b and c. Now consider a different preference over menus, namely the preference $\succ$ such that $x \succ x'$ if and only if $a \in x$ and $a \notin x'$. It is easy to see that this is the preference over menus generated by the Kreps representation (ordinal or additive) when there is only one second stage preference given by the preference we are excluding. In fact, it is not hard to show that this is the *only* such representation possible, so we cannot represent this preference over menus if we exclude the second–stage preference above. In short, any restriction on second stage preferences in Kreps' framework which allows an additive representation of every preference he considers must also allow intrinsically nonadditive ordinal representations of some of the preferences he considers.

[34] It is easy to see that writing u as a monotone transformation of an additive function would not change the argument.

References

[1] F. J. Anscombe and Robert J. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, **34**, 199–205.

[2] Charles Castaing and M. Valadier (1977): *Convex Analysis and Measurable Multifunctions*; Vol. 580 of *Lecture Notes in Mathematics*, Springer Verlag: New York, NY.

[3] Frank H. Clark (1983): *Optimization and Nonsmooth Analysis*; Canadian Mathematical Society series, Wiley: New York, NY.

[4] Eddie Dekel, Barton L. Lipman, and Aldo Rustichini (1997): "Standard State-Space Models Preclude Unawareness," *Econometrica*, forthcoming.

[5] Eddie Dekel, Barton L. Lipman, and Aldo Rustichini (1997): "Recent Developments in Modeling Unforeseen Contingencies," working paper.

[6] Amrita Dhillon and Jean-François Mertens (1996): "Relative Utilitarianism: An Improved Axiomatization," CORE Discussion Paper 9655.

[7] Nelson Dunford and Jacob T. Schwartz (1958): *Linear Operators, Part I: General Theory*, Wiley: New York, NY.

[8] Ron Fagin and Joseph Halpern (1988): "Belief, Awareness and Limited Reasoning," *Artificial Intelligence*, **34**, 39–76.

[9] Peter C. Fishburn (1970): *Utility Theory for Decision Making*; Publications in Operations Research, No. 18, New York: Wiley.

[10] John Geanakoplos (1989): "Game Theory without Partitions, and Applications to Speculation and Consensus," Yale University working paper.

[11] Paolo Ghirardato (1996): "Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty," mimeo, Cal Tech.

[12] Oliver Hart (1995): *Firms, Contracts, and Financial Structure*, Oxford: Clarendon Press.

[13] Edi Karni (1993): "A Definition of Subjective Probabilities with State-Dependent Preferences," *Econometrica*, **61**, 187–198.

[14] Erwin Klein and Anthony C. Thompson (1984): *Theory of Correspondences: Including Application to Mathematical Economics*, Canadian Mathematical Society series, Wiley: New York, NY.

[15] David M. Kreps (1979): "A Representation Theorem for 'Preference for Flexibility'," *Econometrica*, **47**, 565–576.

[16] David M. Kreps (1988): *Notes on the Theory of Choice*, Westview Press: Boulder, CO.

[17] David M. Kreps (1992): "Static Choice and Unforeseen Contingencies" in *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, Partha Dasgupta, Douglas Gale, Oliver Hart, and Eric Maskin (eds.) MIT Press: Cambridge, MA, 259–281.

[18] Eric Maskin and Jean Tirole (1997): "Dynamic Programming, Unforeseen Contingencies, and Incomplete Contracts," mimeo, Harvard University.

[19] Peter Meyer Nieberg (1991): *Banach Lattices*, Springer Verlag: New York, NY.

[20] Salvatore Modica and Aldo Rustichini (1994): "Awareness and Partitional Information Structures," *Theory and Decision*, **37**, 107–124.

[21] Salvatore Modica and Aldo Rustichini (1993): "Unawareness: A Formal Theory of Unforeseen Contingencies: Part II," CORE Discussion Paper.

[22] Klaus Nehring (1996): "Preference for Flexibility and Freedom of Choice in a Savage Framework," mimeo, Cal Davis.

[23] R. Tyrell Rockafellar (1972): *Convex Analysis* , Princeton University Press: Princeton, NJ.

[24] H. L. Royden (1968): *Real Analysis*, Macmillan Publishing Co.: New York, NY.

[25] Leonard J. Savage (1954): *The Foundations of Statistics*, Wiley: New York, NY.

[26] Costis Skiadis (1997): "Subjective Probability under Additive Aggregation of Conditional Preferences," *Journal of Economic Theory*, forthcoming.

[27] Maxwell B. Stinchcombe (1997): "Countably Additive Subjective Probabilities," *Review of Economic Studies*, **64**, 125–146.

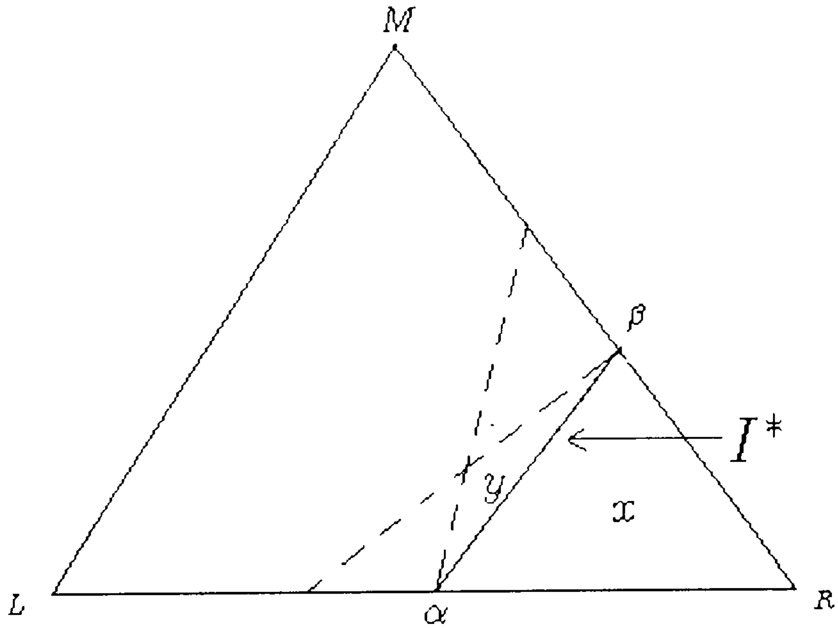

Figure 1