DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

COPING WITH IGNORANCE: UNFORESEEN CONTINGENCIES AND NON-ADDITIVE UNCERTAINTY

Paolo Ghirardato

SOCIAL SCIENCE WORKING PAPER 945

September 1995
Revised May 1996

Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty

Paolo Ghirardato

Abstract

In real-life decision problems, decision makers are never provided with the necessary background structure: the set of states of the world, the outcome space, the set of actions. They have to devise all these by themselves. I model the (static) choice problem of a decision maker (DM) who is aware that her perception of the state space is too coarse, as for instance when there might be unforeseen contingencies. After making assumptions on the way the DM perceives the decision problem, I present a set of axioms on her preferences which imply that they can represented by a (generalized) expectation with respect to a non-additive measure, called a belief function. As it turns out, the very natural axioms presented have strong implications on the way the DM copes with the type of ignorance described above. I show how some decision rules that have been studied in the literature can be obtained as a special case of the model presented here (though they have to be intrepreted differently).

I then show that this formulation of the problem can yield very natural results on the comparative statics of beliefs as the DM's understanding of the decision problem becomes deeper, for instance when unforeseen events become foreseen. I present the implications of these results for a simple asset pricing model. Finally I argue that if we are willing to make assumptions on the faithfulness of the DM's perception with respect to reality, then the DM described here will, in the limit as her perception becomes finer and finer, resemble a Savage DM.

JEL classification numbers: D81, L22

Key words: Unforeseen Contingencies, Belief Functions, Choquet Integrals

Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty

Paolo Ghirardato*

Introduction and Motivation

Few decisions are ever made in the ideal conditions postulated by decision theory. It is assumed that the decision maker (henceforth DM) is *faced* with the decision problem: that is, a set of actions, a set of states of the world and a set of possible outcomes. Moreover in applications to economic problems it is often implicitly assumed that such knowledge of the decision problem is *correct*. That is, the DM knows the "true" way to associate each state of the world and each action with an outcome. But outside the realm of abstract theory, DMs have to *construct* the decision problem before sitting down and making their choices. The set of all options, the set of the states of the world, and the set of the consequences, all have to be thought out and structured. That is probably the hardest task facing a DM who tries to take "good" decisions. Also, there is no reason to believe that such construction will be correct. That is, that it be consistent with the knowledge of a better informed party, who in the following analysis will be called the "modeller".

One of the hardest feats for a DM who is constructing such a decision problem is elaborating the set of all possible state of the world. Each of these is, in Savage's oft-mentioned definition [35, p.9, emphasis added], "a description of the world, leaving no *relevant* aspect undescribed". The word "relevant" holds the key to the problem: given a set of possible actions $\mathcal{F}$, an event A is relevant to the decision problem if at least one

* This paper is a modified version of chapter 1 of my doctoral dissertation at UC Berkeley. I wish to thank my adviser Bob Anderson and Jeff Ely, Larry Epstein, Itzhak Gilboa, Peter Klibanoff, Piero La Mura, Alessandro Lizzeri, Mark Machina, Tom Marschak, Suzanne Scotchmer, Chris Shannon and Daniele Terlizzese for helpful and stimulating discussion. Eddie Dekel, Edi Karni, Debbie Minehart and patient audiences at a variety of Schools and Conferences provided helpful comments on earlier versions. Finally I am grateful to an Editor and two anonymous referees for comments which greatly helped improve the form and substance of the paper. The usual disclaimer applies. Support from an Alfred P. Sloan Doctoral Dissertation Fellowship is gratefully acknowledged. An earlier version was partially funded by Università Bocconi.

action gives a different outcome according to whether A obtains or not. But it seems clear that even for a very simple problem the list of relevant events can be extremely long, making the construction of the state space a superhuman accomplishment. And this is true *a fortiori* if we also imagine that our DM has to make her decision in a very short time.

Similar considerations apply to actions and consequences, of course. However in this work I want to concentrate on the problem of constructing the state space. So I will just assume that the DM can easily think about a list of well-specified actions, and that she can conceive a number of different psychological states in which she could find herself, the possible consequences of the decision problem.[1] When elaborating the space of states of the world, however, the DM thinks that the construction she can come up with is hopelessly crude, in the sense that she is aware of leaving many relevant aspects undescribed.

It is immediate to think about situations in which this will be the case. Take, for instance, the famous political example used by Savage. While it seems sufficiently clear who will be the candidates for the next presidential elections, it is *much* less obvious to say who will be the president of the United States in 2009.[2] And this problem is even more evident in countries with a less stable political system, like most developing countries. For the sake of example, let us consider what would have been the "political" state space for a long term investment by a multinational firm in South Africa to be made in 1990, at a time when Frederick de Clerk had promised that Mandela would be freed from prison and that a new constitution, establishing universal suffrage, would be created. The political scene in South Africa was (and to a large extent still is) far from clear even to insiders. For instance the African National Congress (ANC) was not necessarily the majority party among blacks, and even within the ANC there were (and still are) a variety of factions and currents, some of which have very extreme Marxist-Stalinist positions. Thus, while it is simple to think about states in which ANC dominates the Parliament to be formed, and states in which de Clerk's National Party is ruling, it is not very clear what would really fall inside the catch-all state "none of the above is in power". And the observation about the spectrum of views within the ANC suggests that even the event that ANC is the ruling party leaves a lot of relevant aspects undescribed.

To be a little more concrete, say that the executives of the firm have identified n relevant political events or factors, and let us label E_i the i-th factor. The state space they

[1] The term "psychological states" is Savage's. There is some discussion on whether psychological aspects can be included in the description of consequences in Savage's model, so the reader might prefer to understand consequences as propositions describing purely objective facts, like the bundle of goods that the DM gets to consume. An alternative interpretation of the framework of this model is that objective consequences are well-defined for every state, but their subjective evaluation is not certain *ex ante* (an interpretation suggested by Colin Camerer). This is however not the only interpretation, as you will see below. A very general model which does not employ the notion of consequence is in Skiadas [41] (see the discussion in subsection 2.3).

[2] Former Joint Chief of Staff Colin Powell, who seems to represent a possible candidate for the 2000 elections at the time when this is written, would not even have been considered in 1989, before the Persian Gulf war.

can construct is then given by all the logical products of the E_i and their negation. For example suppose that $n = 2$ and that E_1 is "ANC is ruling party" and E_2 is "(Mandela's) moderate line prevails within ANC'. Then one state is $E_1 \wedge \neg E_2$, the event which is true if and only if ANC takes power and the Mandela line does not prevail (the symbol $\neg$ stands for logical negation, and the symbol $\wedge$ stands for the logical product *and*). There are four of these products. Somebody who has deeper knowledge of the South African situation would be able to think about, say, two more events, E_3 and E_4. Every element of the firm's state space then corresponds to a set of elements in her space: For instance $E_1 \wedge \neg E_2$ will be for her the logical sum of all the states in which $E_1 \wedge \neg E_2$ is true.[3] Let us now associate sets of points with the sets of propositions that we have thus derived. Denote by Π the set corresponding to the firm's state space, and by Ω the set corresponding to the well-informed person's state space. We can now translate our last observation as follows: every element π in the firm's state space is a set of points in her state space. That is, she sees Π as a *partition* of Ω. The important point to observe is that though the executives of the firm are *aware* that Π is too coarse, it is the best they can do given our knowledge and the time they have spent on it, and they have no idea of how it should be refined. (Maybe they could do better if they were given more time to think about it, or they could ask for exeprt advice. But it is clear that, given that decisions have to be taken at some point, coarseness will be the norm rather than the exception.)

Suppose now that the firm were to make some investment f and that $\pi \in \Pi$ obtained. What do the executives expect the outcome $f(\pi)$ to be? They are aware that their construction of the state space is incomplete, so why should they expect the outcome to be determined at all? For example imagine that ANC takes power, but the Mandela line does not prevail. Which type of government, and which policies towards foreign investment, do the executives expect? Given the coarseness in the description of the state the firm might expect a *range* of possible results for its investment.[4] That is, $f(E_1 \wedge \neg E_2)$ is a *set* of outcomes, rather than a singleton. More formally, we might want to be able to think of actions as correspondences from states into outcomes rather than functions, as a way of representing the firm's ignorance. This is the starting point of the analysis in this paper.

The questions I just posed arise especially in modelling DMs facing problems where there are events they are unaware of, commonly known in the literature as "unforeseen

[3] A typical element of her state space is

$$X_1 \wedge X_2 \wedge X_3 \wedge X_4$$

where each X_i is either E_i or $\neg E_i$. Therefore, if $\vee$ is the logical *or* operator, $E_1 \wedge \neg E_2$ formally corresponds to

$$(E_1 \wedge \neg E_2 \wedge E_3 \wedge E_4) \vee (E_1 \wedge \neg E_2 \wedge \neg E_3 \wedge E_4) \vee (E_1 \wedge \neg E_2 \wedge E_3 \wedge \neg E_4) \vee (E_1 \wedge \neg E_2 \wedge \neg E_3 \wedge \neg E_4).$$

[4] This coarseness does not have to be symmetric. For instance it seems quite reasonable to expect a Mandela-led government to pursue very predictable policies (which is what is in fact happening).

contingencies".[5] The issue of unawareness of relevant aspects has been studied recently using the techniques of modal logic, notably by Modica and Rustichini [31, 30] (see also Fagin, Halpern, Moses and Vardi [15]), but the connections between these logical models and the DM's behavior are still an open problem.[6] On the other hand, the problem of modeling choice in the presence of unforeseen contingencies is clearly very important for economic theory. Its relation to the existence of incomplete contracts was suggested by Williamson already in [43]. But while many interesting explanations of incomplete contracts have been offered in recent years (see, e.g., Anderlini and Felli [2], Lipman [27], Aghion and Hermalin [1]), unforeseen contingencies have received much less attention. Probably the only work to date explicitly dealing with this problem is Kreps [26] (originally written in 1986). Kreps's model is axiomatic, as the one in this paper, but the two models are quite different in their assumptions as to how the DM copes with the fact that her state space is under-specified, and which kind of preferences she can elicit (see subsection 2.3 for a detailed discussion). Moreover, while the model to be presented here can be applied to some specific economic issues quite easily (see the short discussion in section 4), applications of his model do not appear to be straightforward.

The objective of this work is to discuss one way in which the DM can cope with her ignorance, and to obtain a mathematical representation for her preferences, in the tradition of decision-theoretic modeling (in particular Savage's fundamental model [35]). It will be assumed that the DM has a *perception* of the decision problem which is given by the state space Π (a partition of the "true" state space Ω) and a correspondence, associating sets of possible outcomes to each π, for each act f. Since I want here to focus on the decision-theoretic implications I will assume that the DM's perception is given, a *datum* of the model. Clearly in doing this I beg one of the fundamental issues: The real challenge lies in explaining how perception comes about. However knowing the decision-theoretic implications of the assumptions we make is a necessary step before proceeding to the important problem of developing a cognitive model. The assumption that Π is a partition of Ω is a rationality requirement: Arguably, some DMs might be completely omitting a subset of the state space. This is what happens when awareness is not, in Modica and Rustichini's term [31], "symmetric" . That is, the DM is aware of the relevance of some event but is not aware of the relevance of its negation. However, as Kreps observes in [26], even in this case it is possible for the DM to make a list of states exhaustive just by creating a catch-all state called "none of the above".[7]

Actions (technically: *acts*) are correspondences from Π to a finite set of consequences

[5] In this work the term "unforeseen contingency" will be used in the very wide sense of any relevant aspect omitted from the description of the state space. The reasoning behind this choice is simple: The behavior of a DM can only be affected by her awareness of her ignorance. In the lack of such awareness, I argue, there is little interesting that a decision theorist (or an economist) can say.

[6] In ongoing research, Modica, Rustichini and Tallon [32] are trying to apply the logical model to a general equilibrium problem, but they approach the problem of modeling the DM's behavior in a way different from that presented in this work.

[7] Another modeling possibility would be assuming that the true state space is a cartesian product of Π and the set of the logical products of all the unforeseen aspects. But this is clearly just a special case of the partition model.

denoted $\mathcal{X}$.[8] In the model no assumption is made on the correctness of this perception of acts: We are only interested in the way the DM sees them. Assuming that the DM has a preference relation on actions, axioms are then given which insure that this relation can be represented mathematically in a certain way. The first six axioms are very similar to Savage's P1-6, with due care being taken of the fact that acts are now correspondences rather than functions. I show that the DM's beliefs on her perceived state space Π can be represented by a (subjective) probability measure P, and then I argue that the axioms imply that the only features of acts which the DM does really care about are the probability distributions that acts induce on the outcome space (suitably defined). The main difference from Savage is that these distributions are not probability measures. They might fail to be *additive*, and associate to the disjoint union of two sets a number which is strictly *larger* than the sum of the numbers associated with the sets taken separately. Technically these distributions are called *belief functions*. The only axiom which does not have a counterpart in Savage's model (the seventh) is a very compelling dominance property which, roughly, says that the DM is not made worse off by more ignorance as long as the latter is related to getting better results.

A crucial feature of the model is a state independence axiom, which as we shall see in this framework embodies an assumption of "complete ignorance". Along with the other axioms, this also implies that the DM "copes with her ignorance" in a simple fashion. This for instance allows the DM's attitude with respect to her ignorance to range from the extreme of being very pessimistic (one who says "given that I don't understand why I might receive x rather than y, I might as well behave as if only the worst outcome y is possible") to the other extreme of being very optimistic (as above, but substitute x with y between quotes).

The mathematical representation of the preferences implied by the axioms has the form of an expectation with respect to the afore-mentioned distributions. However, since belief functions can be non-additive, this expectation is not necessarily a Lebesgue integral. I show however that it is a convex combination of two integrx axioms are very similar to Savage's P1-6, with due care being taken of the fact that acts are now correspondences rather than functions. I show that the DM's beliefs on her perceived state space Π can be represented by a (subjective) probability measure P, and then I argue that the axioms imply that the only features of acts which the DM does really care about are the probability distributions that acts induce on the outcome space (suitably defined). The main difference from Savage is that these distributions are not probability measures. They might fail to be *additive*, and associate to the disjoint union of two

[8] It will be seen later that, as in Savage's model, Π must be an infinite set in this model. This clearly detracts from the interpretation of Π as a set possibly generated by a short list of events. I can justify this as Savage did his model: What is really needed is just that the DM can conceive the existence of an independent randomizing device, like a fair coin. Given that, I only need a finite list of events to create Π, as I can use repeated tosses of the coin to make Π as I rich as I need it for calibration purposes. This is not totally satisfactory, though, as the idea of an independent randomizing device is not philosophically acceptable in a Savage setting. I believe that the model presented here could be reformulated with a finite Π space an infinite consequence space $\mathcal{X}$, following the lines of similar work on Savage's model by Gul [22] and Chew and Karni [6].

sets a number which is strictly *larger* than the sum of the numbers associated with the sets taken separately. Technically these distributions are called *belief functions.* The only axiom which does not have a counterpart in Savage's model (the seventh) is a very compelling dominance property which, roughly, says that the DM is not made worse off by more ignorance as long as the latter is related to getting better results.

A crucial feature of the model is a state independence axiom, which as we shall see in this framework embodies an assumption of "complete ignorance". Along with the other axioms, this also implies that the DM "copes with her ignorance" in a simple fashion. This for instance allows the DM's attitude with respect to her ignorance to range from the extreme of being very pessimistic (one who says "given that I don't understand why I might receive x rather than y, I might as well behave as if only the worst outcome y is possible") to the other extreme of being very optimistic (as above, but substitute x with y between quotes).

The mathematical representation of the preferences implied by the axioms has the form of an expectation with respect to the afore-mentioned distributions. However, since belief functions can be non-additive, this expectation is not necessarily a Lebesgue integral. I show however that it is a convex combination of two integre representation functional can be written a convex combination of Choquet integrals. Section 2 is devoted to a discussion of the model. First (subsection 2.1), I provide axioms for three special cases which give rise to decision rules that are (formally) well-known in the literature on decision theory. Subsection 2.2 is devoted to a discussion of the axioms, focusing in particular on the state independence axiom which plays a key role in the model. Finally subsection 2.3 explains the relation of the present paper to the existing literature. Section 3 discusses the relation between perception and beliefs in more detail. The comparative statics exercise is the object of subsection 3.1, while subsection 3.2 deals with the problem of the relation between truth (i.e., the perception of the modeller) and the DM's perception, a problem novel to the set-up of this model. Section 4 concludes by outlining possible applications and directions for future research.

1 The Model

1.1 Set-up and Axioms

The objects of choice are called *acts.* Each act f is a description of a specific course of action, and the collection of all the descriptions that the DM can think about is denoted $\mathfrak{F}$. Notice that this set contains only acts that are *actually* available to the DM (and she can think of).

The underlying model is described as follows. There is a "true" set Ω of *states of the world*, that is, complete descriptions of (the truth values of) all the events[9] which are

[9] I am using here the word "event" in its naive sense of "a (sufficiently clear) proposition describing

relevant for the decision problem (including events the DM is not aware of, that we could call "unforeseen contingencies"). A typical element of Ω is denoted ω. $\mathcal{X}$ is a *finite* set of *consequences*, possible outcomes of the decision problem. I will assume that $\mathcal{X}$ has n elements, that is, $|\mathcal{X}| = n$. In the underlying model every act $f \in \mathcal{F}$ corresponds to a mapping f_U associating a consequence $f_U(\omega) \in \mathcal{X}$ to every $\omega \in \Omega$, so that $\mathcal{F}$ corresponds to a set of functions from Ω into $\mathcal{X}$.

The DM is not fully aware of the underlying model, which is only known by the modeller. She has a perception of her decision problem which is modeled formally as follows. The set of states of the world that she perceives, denoted Π, is really a class of subsets of Ω, in the sense that every *single* state $\pi \in \Pi$ for the DM is seen by the modeller to be a (possibly large) set of true states. I will assume throughout that the DM constructs Π as an exhaustive list of disjoint events, so that (again from the vantage point of the modeller) it is a *partition* of Ω. I have already given some examples in the Introduction to explain how unawareness (or omission for other causes) of some relevant events will generate such a situation. The DM's understanding of acts is modeled by the *perception correspondence* $\mathbf{V} : \mathcal{F} \times \Pi \to 2^{\mathcal{X}} \setminus \{\emptyset\}$, a mapping which associates to every label $f \in \mathcal{F}$ and every perceived state π a non-empty *set* of possible outcomes,[10] to be interpreted as all the outcomes considered by the DM to be possible results of f in state π. A simplification of the notation is obtained by writing $f_V(\pi)$ instead of $\mathbf{V}(f, \pi)$, so that $\mathbf{V}$ can also be seen as a rule associating a *correspondence* $f_V : \Pi \to 2^{\mathcal{X}} \setminus \{\emptyset\}$ to the label f. The pair $< \Pi, \mathbf{V} >$ forms what I will call the *epistemic status* of the DM. Except where otherwise noted, it will be fixed and given.

It should be stressed that formally there need not be any relation between $f_V(\pi)$, the perceived outcomes of f in state π, and $f_U(\pi)$, the actual set of outcomes. That is, the DM's perception could be totally "wrong". Correctness of perception plays an important role only in sequential applications, i.e., situations in which the choice problem is faced more than once. What is really important for the particular nature of this model is that we are allowing the DM to cope with her ignorance (as expressed by the fact that Π is a partition of something much finer, and the DM is aware of this) by letting her construct acts as correspondences rather than functions. The latter aspect is what distinguishes formally this model from Savage's classical SEU model [35].

Remark 1 As only the DM's epistemic status matters when modelling her preferences, except where otherwise indicated I will abuse notation and delete the $_V$ subscript from f_V, thus identifying the label f with its formal representation in the mind of the DM. Thus $\mathcal{F}$ will now be a set of correspondences from Π into $\mathcal{X}$. Once the subjective nature of the representation is well understood, this should not create any confusion. ◇

An act $f \in \mathcal{F}$ which is perceived as a function (that is, $f(\pi)$ is a singleton for every

some happening of the world, which can be verified to be true or false at some point in the future", not to mean a set of points. A rigorous definition of the notion of state of the world, using modal logic, is possible but beyond the scope of this paper.

[10] For any set A, 2^A denotes its *power set*, i.e., the set of all its subsets.

$\pi \in \Pi$) is said to be *crisp*.[11] It will be shortly clear that if all acts in $\mathcal{F}$ are crisp then the DM believes she is facing a standard problem, and what follows will reduce to Savage's model.

The DM's preferences are formalized by a binary relation $\succeq$. As was the case for Savage, I need $\succeq$ to be defined on a set larger than $\mathcal{F}$ so as to be able to obtain enough structure for the representation. The definition of this domain is in this model somewhat complicated by the fact that I do not want to put arbitrary restrictions on the DM's perception of her ignorance. The following example illustrates the problem.

Example 1 Think again about the problem of the multinational investing in South Africa in 1990. The problem which the executives face is, as we mentioned, writing a sufficiently detailed state space of the political situation in South Africa during the life of the investment. In particular, it is not very clear what will happen if in the 1994 elections the National Party (NP) loses the elections (an event I'll label E_1): Obviously a crude sub-state space would be given by the event E_2, that the ANC wins the elections, and its negation. But how the state space should be further subdivided is quite unclear, especially if $\neg E_2$. So conditionally on $\neg E_1$, the executives might think that every investment is connected with a high degree of ignorance, and the more so if $\neg E_2$. This would be reflected by $f(\pi)$ being possibly very large for every investment f and every substate π of the event $\neg E_1$. On the other hand conditionally on E_1 happening, the executives might be less concerned with unpredictable events. They might believe that if the NP wins a lawfully held election, then the political situation in S.Africa will be quite stable, so that the outcomes of an investment will be quite predictable (of course they might still be uncertain, but they depend on factors that the firm is well aware of). So conditionally on E_1 it might be that for every f, $f(\pi)$ is either a singleton or quite small: in our terminology, acts are closer to being crisp. $\triangle$

The immediate extension of Savage's [35] model is to let the domain of preferences be the set of *all* the correspondences from Π into $\mathcal{X}$. But doing so is equivalent to asking the executives of the firm in the example to consider the counterfactual situation in which the political situation in S.Africa is stable, since NP won the elections, and yet the executives are afraid of high instability. While it is possible that they might be willing to entertain similar doubts, it is also possible that they would consider such possibility farfetched, so that their preferences lose meaning. The point is that defining preferences on the large domain mentioned above might involve asking the DM to consider herself in a counterfactual epistemic situation (i.e., a situation in which she feels more ignorant than she is). The approach taken below, while requiring some more work in terms of definitions, has the advantage of avoiding to ask the DM this type of counterfactual questions. As we will see shortly, it *generalizes* the Savage-like approach mentioned above, since it yields the same results when the DM feels "uniformly ignorant" on Π.

[11] In what follows I will be using the natural identification of $x \in \mathcal{X}$ with the singleton subset $\{x\} \in 2^{\mathcal{X}}$.

Partition the set $2^{\mathcal{X}}$ as follows. For $i = 1, 2, \ldots, n$, let

$$X(i) \equiv \{Y \in 2^{\mathcal{X}} : |Y| = i\}.$$

That is, $X(i)$ is the class of all sets of i different consequences. Clearly $\{X(i)\}_i$ is a partition of $2^{\mathcal{X}}$. Given $f \in \mathcal{F}$, for every $X \subseteq \mathcal{X}$ let $f^{-1}(X) \equiv \{\pi \in \Pi : f(\pi) = X\}$ and $f^{-1}(X(i)) \equiv \cup_{X \in X(i)} f^{-1}(X)$. Then $\{f^{-1}(X(i))\}_i$ is a partition of Π. Finally, for $i = 0, \ldots, n-1$, define

$$I_{n-i} \equiv \left[\bigcup_{f \in \mathcal{F}} f^{-1}(X(n-i)) \right] \setminus \left(\bigcup_{j=0}^{i-1} I_{n-j} \right). \tag{1}$$

Then $\{I_i\}_i$ is a partition of Π with the property that the DM considers it possible to obtain (i.e., there is an act $f \in \mathcal{F}$ which delivers) up to n different results in all states in I_n, up to $n-1$ on I_{n-1}, etc. Notice that this is constructed using only $\mathcal{F}$, the set of (subjective perceptions of) acts that the DM considers factually possible.

Now we can present the set of acts that the DM will be asked to rank. It should be stressed that I do not assume that these acts are available to the DM, but just that she can express a preference over them. Let

$$\bar{\mathcal{F}} \equiv \{f : \Pi \to 2^{\mathcal{X}} : \forall i, \forall \pi \in I_i, \ |f(\pi)| \leq i\} \tag{2}$$

Thus we ask the DM to evaluate only correspondences which are "feasible" in the sense of embodying possible levels of ignorance. As I mentioned earlier, the set $\bar{\mathcal{F}}$ will in general be smaller than the set $(2^{\mathcal{X}})^{\Pi}$ of all correspondences from Π into $\mathcal{X}$. It will be equal to it in the case in which $\Pi = I_n$, that I will define the "uniform ignorance" case. Notice that if $\mathcal{F}$ contains only crisp acts, then $\bar{\mathcal{F}} = \mathcal{X}^{\Pi}$, the set of all *functions* from Π into $\mathcal{X}$. Hence the set of acts will just be the one used by Savage, and the DM described here will conform to the SEU model. The subset of all the crisp acts in $\bar{\mathcal{F}}$ will be denoted $\bar{\mathcal{F}}_C$.

It should be stressed that assuming that the DM has (complete) preferences on $\bar{\mathcal{F}}$ is not free of the problem of asking her to rank possibly counterfactual acts. In fact observe that, as in Savage, $\bar{\mathcal{F}}$ contains all "constant" acts of the form $f \equiv x$, for $x \in \mathcal{X}$. Clearly most of these are purely theoretical, but it is by asking the DM to rank all constant acts that we obtain her preference relation on consequences. That is we say that $x \succeq y$, for $x, y \in \mathcal{X}$, if $f \succeq g$, where $f \equiv x$ and $g \equiv y$. Notice that constant acts are crisp. While I believe that we could under quite general circumstances do without asking this type of questions,[12] here in the sake of clarity and simplicity I will follow the traditional route.

One might wonder at this point if a similar exercise can be done to compare any two *subsets* of outcomes $X, Y \subseteq \mathcal{X}$. In the uniform ignorance case this would be done exactly in the same way. That is, say that $X \succeq Y$ if $f \succeq g$, where $f \equiv X$ and $g \equiv Y$, since the

[12] This would proceed in a fashion analogous to the one I will use below to define a preference on sets. A key assumption for the method to work is that there is a non-null event conditionally on which the DM can conceive receiving every possible consequence.

latter then belong to $\bar{\mathfrak{F}}$. In general an ordering between subsets can be established as follows. Given $X \in 2^{\mathfrak{X}}$, suppose that $|X| = i$ and let $I(X) \equiv \cup_{h \geq i} I_h$. $I(X) \subseteq \Pi$ is the largest event on which X can be the result of an act in $\bar{\mathfrak{F}}$. So for a pair $X, Y \in 2^{\mathfrak{X}}$ let $I = I(X) \cap I(Y)$. Then *if* $I \neq \emptyset$ define, for some $x \in \mathfrak{X}$,

$$X \succeq Y \iff (X, I; x, I^c) \succeq (Y, I; x, I^c) \tag{3}$$

where $(Z, I; x, I^c)$ is the act which yields the set of outcomes $Z \in \{X, Y\}$ for every $\pi \in I$ and x otherwise.[13] Axiom 2 will imply that the choice of x is irrelevant. What definition (3) does is to ask the DM to compare sets of outcomes only where her perception of ignorance would theoretically allow it. (Notice that both acts belong to $\bar{\mathfrak{F}}$.)[14]

Remark 2 It should be observed that the ordering on $2^{\mathfrak{X}}$ thus constructed is not necessarily complete (and this even if the preference on $\bar{\mathfrak{F}}$ is). In fact if $I_h = \emptyset$ for all $h \geq k$ then X and Y cannot be compared. For instance, a DM which is described by Savage's model will not be able to compare any two non-singleton sets. ◇

$\succeq$ is assumed to satisfy the following axioms. As usual I will denote $\succ$ and $\sim$ respectively the asymmetric and symmetric part of $\succeq$.

Axiom 1 (Weak order) $\succeq$ *is transitive and complete on* $\bar{\mathfrak{F}}$.

Axiom 2 (Sure-thing principle) *If* $f, g, h, h' \in \bar{\mathfrak{F}}$, $A \subseteq \Pi$, *then*

$$\begin{bmatrix} f, & A \\ h, & A^c \end{bmatrix} \succeq \begin{bmatrix} g, & A \\ h, & A^c \end{bmatrix} \iff \begin{bmatrix} f, & A \\ h', & A^c \end{bmatrix} \succeq \begin{bmatrix} g, & A \\ h', & A^c \end{bmatrix}.$$

These axioms are well-known. In calling Axiom 2 "sure-thing principle" I am faithful to the literature in Decision Theory, but unfaithful to Savage and Fishburn [16], who give this name to Savage's axiom P7. Notice that axiom 2 requires the sure-thing principle to hold only for events that the DM can identify (subsets of Π), not for all events in the underlying state space (subsets of Ω).[15] It is formally identical to Jaffray and Wakker's "weak sure-thing principle" [25]. Given axioms 1 and 2 we can define "conditional preferences" as follows: for every $A \subseteq \Pi$ let $f \succeq_A g$ if $(f, A; h, A^c) \succeq (g, A; h, A^c)$ for some (hence all, by axiom 2) $h \in \bar{\mathfrak{F}}$. One can easily see that all $\succeq_A$ are complete and transitive.

Definition 1 $A \subseteq \Pi$ *is null if* $f \sim_A g$ *for every* $f, g \in \bar{\mathfrak{F}}$.

[13] In general: For $A \subseteq \Pi$, A^c denotes the complement of A in the DM's knowledge (that is, $A^c \equiv \{\pi \in \Pi : \pi \notin A\}$), and $(f, A; g, A^c)$ denotes the act which coincides with f on A and with g on A^c.

[14] Of course this again requires the DM to rank acts which might be counterfactual in their results, very much analogously to what is done in the Savage model. But the definition of $\{I_i\}$ implies that in contructing this relation on $2^{\mathfrak{X}}$ we never ask the DM questions which are epistemically counterfactual.

[15] I am abusing notation a bit here. Formally A is a set of subsets of Ω, and not a subset itself. In order not to encumber notation I also use A to mean the union (as seen by the modeller) of all the subsets belonging to A.

Axiom 3 (State independence/Ignorance) *If $A \subseteq \Pi$ is not null and if, given $X, Y \subseteq \mathfrak{X}$, there are $f, g \in \bar{\mathfrak{F}}$ such that $f(\pi) = X$ and $g(\pi) = Y$ for all $\pi \in A$, then*

$$f \succeq_A g \Longleftrightarrow X \succeq Y.$$

This axiom is the natural extension of Savage's P3 to the case in which acts are correspondences. It extends the state independence of P3 to the comparison of sets of outcomes. Observe that the clause that f and g belong to $\bar{\mathfrak{F}}$ implies that X and Y are comparable. It is unnecessary if X and Y are singletons. I will discuss axiom 3 in detail in subsection 2.2, after presenting the representation theorem. As will be seen more clearly there, this axiom embodies a complete ignorance assumption: The evaluation of a set of results X (which is more than a singleton because of the DM's awareness of her ignorance) does not depend on the circumstances in which it is obtained (i.e., the event A and the act f which yields it), which is a way of saying that the DM really *does not have a clue* as to what justifies different outcomes within X.

Axiom 4 (Comparative probability) *Suppose that $X_1, X_2, X_3, X_4 \in 2^{\mathfrak{X}}$ are such that $C = \cap_{i=1}^4 I(X_i) \neq \emptyset$. Then for every $A, B \subseteq C$ and every $f \in \bar{\mathfrak{F}}$, $X_1 \succ X_2$ and $X_3 \succ X_4$ imply*

$$\begin{bmatrix} X_1, & A \\ X_2, & C \setminus A \\ f, & C^c \end{bmatrix} \succeq \begin{bmatrix} X_1, & B \\ X_2, & C \setminus B \\ f, & C^c \end{bmatrix} \Longleftrightarrow \begin{bmatrix} X_3, & A \\ X_4, & C \setminus A \\ f, & C^c \end{bmatrix} \succeq \begin{bmatrix} X_3, & B \\ X_4, & C \setminus B \\ f, & C^c \end{bmatrix}. \tag{4}$$

Axiom 5 (Non-triviality) *There are $x, y \in \mathfrak{X}$ such that $x \succ y$.*

Axiom 5 is a standard non-triviality assumption. It is not very restrictive, and it is synctactically identical to Savage's P5. Regarding axiom 4, observe that the definition of C implies that all the acts in (4) belong to $\bar{\mathfrak{F}}$. Analogously to Savage's P4, axiom 4 insures that we can use some crisp acts (*bets*) to construct a binary relation on 2^{Π} as follows: For $A, B \subseteq \Pi$, $x, y \in \mathfrak{X}$ such that $x \succ y$ (which exist by axiom 5), define

$$A \succeq^* B \Longleftrightarrow \begin{bmatrix} x, & A \\ y, & A^c \end{bmatrix} \succeq \begin{bmatrix} x, & B \\ y, & B^c \end{bmatrix}. \tag{5}$$

By axiom 4 this definition does not depend on the choice of the singletons x and y (notice that in this case $C = \Pi$). Notice that both the acts in the right-hand side of (5) are crisp. We call $\succeq^*$ a *likelihood* relation and if (5) holds we say that A is *at least as likely as* B. It is easy to verify that if all acts in $\mathfrak{F}$ (hence in $\bar{\mathfrak{F}}$) are crisp axiom 4 is equivalent (in the presence of axiom 2, which is then equivalent to P2) to Savage's P4.[16] In the uniform ignorance case (in which $C = \Pi$ regardless the choice of subsets in the axiom), it is equivalent to the natural extension of P4 in which lowercase letters are substituted with uppercase letters.

[16] An axiom which lies halfway between P4 and axiom 4 was used by Epstein and Le Breton [12] in a different context (in which P2 was not assumed) to insure that it is possible to define conditional likelihood relations in a fashion similar to (5).

Axiom 6 (Archimedean) *If $f, g \in \bar{\mathcal{F}}$ are such that $f \succ g$ and $x \in \mathcal{X}$ then there is a finite partition $\mathcal{H}$ of Π such that, for every $H \in \mathcal{H}$:*

(i) $(x, H; f, H^c) \succ g$,

(ii) $f \succ (x, H; g, H^c)$.

This is similar to Savage's P6. It imposes a continuity property on preferences, since it says that for every *singleton* consequence x, we can partition Π finely enough so that substituting x for the result of some act on an element of the partition does not change preferences. However, it is weaker than the natural counterpart of P6 in the case in which acts are correspondences rather than functions, since it only requires the axiom to hold for a singleton x rather than a set of results. Like P6, it implies that Π must at least be countable.[17]

Axioms 1-6 immediately imply that the DM's beliefs can be represented by a probability measure on Π.

Theorem 1 (Savage, 1954) *If axioms 1-6 hold then $\succeq^*$ can be represented by a unique convex-ranged probability measure on Π. That is, there is a finitely additive probability P, defined on 2^Π, such that for every $A, B \subseteq \Pi$*

$$A \succeq^* B \Longleftrightarrow P(A) \geq P(B) \tag{6}$$

and, if $A \subseteq \Pi$ and $0 \leq \rho \leq 1$, there is $B \subseteq A$ such that

$$P(B) = \rho P(A). \tag{7}$$

Proof: Just notice that axioms 1-6 imply that P1-6 of Savage hold on $\bar{\mathcal{F}}_C$ and apply Savage's result (see Fishburn [16], sections 14.2 and 14.3). ∎

In the standard Savage framework, every act f induces via P a probability measure P_f on $\mathcal{X}$, called the *distribution induced by* f, as follows, for every $X \subseteq \mathcal{X}$,

$$P_f(X) \equiv P(\{\pi \in \Pi : f(\pi) \in X\}). \tag{8}$$

In our set-up, since $f \in \bar{\mathcal{F}}$ is not necessarily crisp, the definition above can be meaningless. A natural extension to correspondences is the following,[18] for every $X \subseteq \mathcal{X}$, let

$$\nu_f(X) \equiv P(\{\pi \in \Pi : f(\pi) \subseteq X\}). \tag{9}$$

Notice that $\nu_f = P_f$ when f is crisp, so (9) is a proper generalization of (8). When f is not crisp ν_f satisfies most of the properties of a probability measure, but not all. In particular it is not additive. An exact statement of its properties requires a short aside.

[17] Savage also had another axiom, P7, which we shall not need since I took the consequence space to be finite (so that all our acts will be "simple" in the sense of Savage).

[18] One could argue that it is no less natural to use instead $\bar{\nu}_f$ as defined in (21) below. However (19) shows that ν_f uniquely defines $\bar{\nu}_f$ and *vice versa*. This immediately implies that the choice of $\bar{\nu}_f$ or ν_f is non-consequential for the analysis to follow.

An Aside: Capacities and Belief Functions

A (normalized) *capacity* on $\mathcal{X}$ is a set-function $\nu : 2^{\mathcal{X}} \to \mathbf{R}$ which satisfies the following properties:

(i) $\nu(\emptyset) = 0$, $\nu(\mathcal{X}) = 1$,

(ii) $\forall A, B \in 2^{\mathcal{X}} : A \subseteq B \Rightarrow \nu(A) \leq \nu(B)$.

Thus a capacity is monotonic, but it is not necessarily additive. A *belief function* (Shafer [38]) is a capacity which also satisfies the following (unintuitive) property, called *total monotonicity*:

(iii) for every $n > 0$ and every collection $A_1, \ldots, A_n \in 2^{\mathcal{X}}$

$$\nu(\cup_{i=1}^n A_i) \geq \sum_{\substack{I \subset \{1,\ldots,n\} \\ I \neq \emptyset}} (-1)^{|I|+1} \nu(\cap_{i \in I} A_i)$$

where $|I|$ is the cardinality of set I.

Probability measures are belief functions (by the well-known *inclusion-exclusion formula*), but the converse is not true. In particular notice that belief functions are *supermodular* (or *convex*) capacities, i.e., for every $A, B \in 2^{\mathcal{X}}$ we have $\nu(A \cup B) + \nu(A \cap B) \geq \nu(A) + \nu(B)$.

Admittedly, property (iii) is quite obscure. An equivalent characterization of belief functions which might be more palatable is the following. Given a probability measure φ on the power set[19] $2^{\mathcal{X}}$, define a capacity ν on $\mathcal{X}$ as follows

$$\nu(X) = \sum_{A \subseteq X} \varphi(A). \tag{10}$$

Then it is possible to show that ν is a belief function (see Shafer [38]). Conversely, for every belief function ν on $\mathcal{X}$ there is a probability measure φ on $2^{\mathcal{X}}$ which satisfies (10), called its *Möbius transform*. The standard interpretation for $\varphi(X)$ is that it is the amount of likelihood weight which is *specifically* assigned to event $X \subseteq \mathcal{X}$ (i.e., it cannot be divided among its subsets), and from (10) we get that $\nu(X)$ represents the total likelihood of the occurrence of X. ◇

Now, consider the set-function on $\mathcal{X}$ thus defined: for $X \subseteq \mathcal{X}$ let

$$\varphi_f(X) \equiv P(\{\pi \in \Pi : f(\pi) = X\}). \tag{11}$$

It is obvious that φ_f is extendable by additivity to a probability measure on $2^{\mathcal{X}}$, and it satisfies (10) with ν replaced by ν_f as defined in (9). This shows that the distribution ν_f

[19] Formally, $\varphi : 2^{2^{\mathcal{X}}} \to [0,1]$, that is, φ assigns a weight to every *family* of subsets of $\mathcal{X}$. I will abuse notation by writing $\varphi(A)$ for $A \in 2^{\mathcal{X}}$ in place of the more rigorous $\varphi(\{A\})$.

is a belief function, and φ_f is its Möbius transform. One immediately checks that when f is crisp, φ_f assigns zero weight to all non-singleton subsets of $\mathfrak{X}$, which again yields $\nu_f = P_f$.

The next step is to ask whether the distribution ν_f plays any special role in the formal description of the DM's preferences. In the case where all acts in $\mathfrak{F}$ are crisp, Savage proved that if a DM's preference relation satisfies axioms 1-6 then she is indifferent between any two acts which induce the same distribution. So the distribution over $\mathfrak{X}$ is the *only* feature of an act which the DM cares for (what Kreps calls the "axiom zero" of the von Neumann-Morgenstern model). Clearly the same result applies here to comparisons of acts in $\bar{\mathfrak{F}}_C$. In fact this is true for all acts in $\bar{\mathfrak{F}}$, as the next proposition shows.[20]

Proposition 1 *Suppose that $\succeq$ satisfies axioms 1-6 and $f, g \in \bar{\mathfrak{F}}$ are such that $\nu_f = \nu_g$. Then $f \sim g$.*

So we see that axioms 1-6 imply that in a sense the DM's "beliefs" are, in the framework of this model, naturally non-additive. Careful analysis of the proof reveals that the main responsibility for this result is of axioms 3 and 4. Axiom 4 plays a role in insuring that the evaluation of probabilities of events in Π does not depend on the acts which we are considering (be they crisp or not), and axiom 3 implies that the evaluation of sets of outcomes does not depend on either the acts which yield them, nor the circumstances under which they are obtained (see subsection 2.2 for more discussion).

Given axioms 1-6, it would now be possible to follow the steps of Savage's proof of his theorem to show that the DM's preferences have a subjective expected utility representation of the following type: There is a function $U : 2^{\mathfrak{X}} \to \mathbf{R}$ such that, for every $f, g \in \bar{\mathfrak{F}}$,

$$f \succeq g \Longleftrightarrow \int_\Pi U(f(\pi))\, dP(\pi) \geq \int_\Pi U(g(\pi))\, dP(\pi).$$

But this representation is incomplete in an important respect: It does not characterize completely the DM's attitude with respect to her "ignorance". That is, suppose that $X = \{x_1, \ldots, x_m\}$ is the possible result of some act in some circumstances. How is the DM going to compare X with its proper subsets? For instance, would she be happier if she could choose an act which under the same circumstances yields the singleton x_1? The axioms we have seen so far imply that such comparison will not depend on the act in consideration, or the circumstances, but they go no further than that. The following axiom, by restricting these comparison in a fairly intuitive way, surpirisingly provides all the structure we need.

Axiom 7 (Dominance) *Given $X, Y \in 2^{\mathfrak{X}}$, suppose that X and Y are comparable (i.e., $I(X) \cap I(Y) \neq \emptyset$) and that for every $x \in X$, $x \succeq y$ for all $y \in Y$. Then $X \succeq Y$.*

[20] The proof of this and the following results are found in the appendix. Jaffray and Wakker [25], have this as an axiom, which they call the "neutrality axiom".

As it is apparent, this is a rather weak dominance assumption, and it does not appear to be very restrictive: Given that all results in X are better than all results in Y, even if the DM feels considerably ignorant in both cases, she will prefer "receiving" X to Y. It is now possible to prove the following interesting result.

Proposition 2 *Suppose that $A, B \subseteq \Pi$, A non-null, $f, g \in \bar{\mathfrak{F}}$ are such that for some $X \subseteq \mathfrak{X}$, $A \subseteq f^{-1}(X)$, $B \subseteq g^{-1}(X)$ and $A \sim^* B$. If $X = \{x_1, \ldots, x_m\}$ relabel so that $x_1 \succeq x_2 \succeq \cdots \succeq x_m$. Then if $\succeq$ satisfies axioms 1-7 there are $C \subseteq A$ and $D \subseteq B$ such that $C \sim^* D$ and*

$$f \sim (x_m, C;\ x_1, A \setminus C;\ f, A^c), \tag{12}$$
$$g \sim (x_m, D;\ x_1, B \setminus D;\ g, B^c). \tag{13}$$

This proposition really contains two results. Suppose that an act f yields a set of consequences X for all states in some non-null event A. Then first of all it is shown that we can find a $C \subseteq A$ such that *conditionally on* A the DM is indifferent between f and the "mixture" act which yields the worst outcome in X over C and the best outcome in X over $A \setminus C$. In fact notice that, using the definition of conditional preferences (12) can be rewritten as

$$f \sim_A (x_m, C;\ x_1, A \setminus C; f, A^c). \tag{14}$$

This is a continuity property of preferences, and it is mostly the work of axiom 6. More importantly, proposition 2 also proves that such substitution can be carried out in a consistent manner, in the sense that given another act g which yields the same X on a set B of equal probability, the set $D \subseteq B$ for which

$$g \sim_B (x_m, D;\ x_1, B \setminus D; g, B^c), \tag{15}$$

must have the same probability as C. So what matters in substituting a set X of consequences with a "mixture" of its best and worst elements is just the identity of X, not under which circumstances, or as a result of which act, X is received. This is mostly due to the state independence of preferences implied by axioms 3 and 4. Theorem 4 in the appendix shows that the result in proposition 2 extends to the case in which A and B have different (positive) probability: That is, there are C and D for which (12) and (13) hold such that $P(C|A) = P(D|B)$, where the conditional probabilities are defined in the usual way (eq. (36) in the appendix). This latter conclusion is crucial in characterizing the representation.

An immediate consequence of proposition 2 is that for every non-crisp act $f \in \bar{\mathfrak{F}}$ it is possible to find a crisp act $f^* \in \bar{\mathfrak{F}}_C$ which is indifferent to it, a result which is the main ingredient for the representation theorem.

Corollary 1 *If $\succeq$ satisfies axioms 1-7 then for every $f \in \bar{\mathfrak{F}}$ there is $f^* \in \bar{\mathfrak{F}}_C$ such that $f \sim f^*$.*

We have thus set the stage for the leading character of this story.

1.2 The Representation Theorem

The main result of the paper is

Theorem 2 *$\succeq$ satisfies axioms 1-7 if and only if there is a convex-ranged probability measure P on Π, a non-constant utility function $u : \mathfrak{X} \to \mathbf{R}$ and a function $\alpha : 2^{\mathfrak{X}} \to [0,1]$ such that, for every $f, g \in \mathfrak{F}$,*

$$f \succeq g \Longleftrightarrow U(f) \geq U(g), \tag{16}$$

where U is defined as follows:

$$U(f) \equiv \sum_{X \subseteq \mathfrak{X}} \varphi_f(X) \left[\alpha(X) \min_{x \in X} u(x) + (1 - \alpha(X)) \max_{x \in X} u(x)\right]. \tag{17}$$

φ_f is defined as in (11). P and α are unique and u is unique up to a positive affine transformation.

The proof of the theorem can be found in the appendix, but its structure at this point should be clear. We have already seen in theorem 1 how to derive the measure P. Corollary 1 told us that every act is indifferent to some crisp act. Thus we can apply Savage's representation theorem to the set of crisp acts, to obtain a subjective expected utility representation on $\bar{\mathfrak{F}}_C$, and use that to represent preferences on $\bar{\mathfrak{F}}$. It is then a simple matter to show the representation has the form given in equation (17), given the specific way in which the indifferent crisp act is constructed. Proposition 2 (more correctly: its extension in theorem 4) is used to show that the weight of the mixture between the worst and best element of X depends only on X, so that it can be written as $\alpha(X)$.

As the expert reader surely noticed, the representation has a familiar aspect: It reminds of Hurwicz's pessimism-optimism index model (see Arrow and Hurwicz [4]). In fact the axiomatization presented here provides a justification for such decision model in a more general situation than the "complete ignorance" studied by Arrow and Hurwicz (followed in this by Jaffray and Wakker [25]).[21] This is why I call α the *pessimism index function*. That this is so can be seen from the mathematical representation, as I will discuss in the next subsection. But it can also be immediately seen by considering how it is obtained. In fact, using the notation of proposition 2, $\alpha(X) = P(C|A)$. If the DM is very pessimistic with regards to her own ignorance, she will behave as if X is only slightly better than its worst element. In other words she will have a "large" C in equation (12), which gives an $\alpha(X)$ close to 1.[22]

[21] Moreover, I argue that this paper suggests a more convincing justification for the degree of "complete ignorance" which it allows the DM to have: What better reason is there for being completely ignorant on some relevant aspects of the state space than being unaware of them?

[22] Clearly this intuition is valid only insofar as receiving a nonsingleton set of outcomes X makes sense for the DM. In other words, a high $\alpha(X)$ represents a pessimistic attitude with respect to ignorance *only if* some ignorance is perceived.

1.3 The Representation is Nothing but...

It turns out that it is possible to rewrite the representation (16) in a way which is very familiar to decision theorists. Unfortunately this requires introducing some notation and preliminary results, which call for another aside.

An Aside: Choquet Integrals

Suppose that S is some space and Σ is an algebra of subsets of S. An interesting question is: how do we integrate with respect to a capacity (not necessarily a belief function) defined on Σ? A concept of integral for capacities was proposed by Choquet in [7]. If $a : S \to \mathbf{R}$ is a bounded Σ-measurable function and ν is a capacity on S we define the *Choquet integral* of a with respect to ν to be the number

$$\begin{aligned}\int_S a(s)\, d\nu(s) &= \int_0^\infty \nu(\{s \in S : a(s) \geq \alpha\})\, d\alpha \\ &\quad + \int_{-\infty}^0 [\nu(\{s \in S : a(s) \geq \alpha\}) - 1]\, d\alpha\end{aligned}$$

where the integrals are taken in the sense of Riemann. In particular, if S is finite with cardinality n and $a(s_1) \geq a(s_2) \geq \ldots \geq a(s_n)$, then

$$\int_S a(s)\, d\nu(s) = \sum_{i=1}^{n-1} (a(s_i) - a(s_{i+1}))\nu(\{s_1, \ldots, s_i\}) + a(s_n). \tag{18}$$

Note that since the integrands are monotone, the Choquet integral always exists. Also if ν is a probability the integral reduces to a standard Lebesgue integral. Suppose now that S is finite and that $\Sigma = 2^S$. For every capacity ν on S we can obtain its Möbius transform φ using the following formula:[23] for every $A \subseteq S$,

$$\varphi(A) = \sum_{B \subseteq A} (-1)^{|A \setminus B|}\, \nu(B).$$

Given a supermodular capacity ν on S define its *complementary* capacity $\bar{\nu}$ as follows, for every $A \subseteq S$,

$$\bar{\nu}(A) = 1 - \nu(A^c). \tag{19}$$

A well-known result (see, e.g., Gilboa and Schmeidler [21, theorem 4.3]) shows that for every function $a : S \to \mathbf{R}$ the Choquet integral with respect to a capacity nu can be written as

$$\int_S a(s)\, d\nu(s) = \sum_{A \subseteq S} \varphi(A)\, [\min_{s \in A} a(s)],$$

[23] It is easy to see, *mutatis mutandis*, that ν_f and φ_f as defined above satisfy this by construction.

and simmetrically

$$\int_S a(s)\, d\bar{\nu}(s) = \sum_{A \subseteq S} \varphi(A)\, [\max_{s \in A} a(s)].$$

Given a capacity ν on (S, Σ) define the *core* of ν, and label $C(\nu)$, the set of all the (finitely additive) probability measures which dominate ν, i.e.,

$$C(\nu) \equiv \{P : P \text{ is a probability on } (S, \Sigma),\ \forall A \in \Sigma,\ P(A) \geq \nu(A)\}. \tag{20}$$

Notice that the $C(\nu)$ is also the set of the probabilities which are dominated by $\bar{\nu}$, as

$$P(A) = 1 - P(A^c) \leq 1 - \nu(A^c) = \bar{\nu}(A).$$

In general the core of a capacity might be empty, but Shapley proved [39] that the core of a supermodular capacity is non-empty, and Schmeidler showed [36, Proposition] that for every bounded Σ-measurable function $a : S \to \mathbf{R}$, ν is a supermodular capacity if and only if

$$\int_S a(s)\, d\nu(s) = \min_{P \in C(\nu)} \int_S a(s)\, dP(s),$$

or, equivalently, if and only if

$$\int_S a(s)\, d\bar{\nu}(s) = \max_{P \in C(\nu)} \int_S a(s)\, dP(s).$$

This result tells us that we can interpret a supermodular capacity as a "sufficient statistic" for a class of measures. The Choquet integral with respect to such capacity is a functional which associates with every function a the (standard) integral with respect to the "worst" possible measure in the set. Dual results hold for submodular capacities (like $\bar{\nu}$), for then the Choquet integral integrates a with respect to the "best" measure in the set. In this sense a Choquet integral with respect to a supermodular (resp. submodular) capacity represents an extremely pessimistic (resp. optimistic) attitude. $\diamond$

Notice that the distribution ν_f defined in (9) is supermodular (being a belief function) and that its complementary capacity $\bar{\nu}_f$ is equal to, for all $X \subseteq \mathfrak{X}$,

$$\bar{\nu}_f(X) = P(\{\pi \in \Pi : f(\pi) \cap X \neq \emptyset\}). \tag{21}$$

For a DM whose preferences can be represented as in (16), let $w_f(\alpha) = \sum_{X \subseteq \mathfrak{X}} \varphi_f(X)\alpha(X)$. For all $X \subseteq \mathfrak{X}$, let

$$\chi_f(X) = \frac{\varphi_f(X)\alpha(X)}{w_f(\alpha)}, \quad \psi_f(X) = \frac{\varphi_f(X)(1 - \alpha(X))}{1 - w_f(\alpha)}.$$

Observe that χ_f and ψ_f satisfy all the required properties of a Möbius transform, so that they induce via (10) two belief functions, respectively denoted ν_f^α and μ_f^α. It is now immediate to check that we can rewrite the preference functional $U(f)$ in eq. (17) as follows:

$$U(f) = w_f(\alpha) \sum_{X \subseteq \mathcal{X}} \chi_f(X) \min_{x \in X} u(x) + (1 - w_f(\alpha)) \sum_{X \subseteq \mathcal{X}} \psi_f(X) \max_{x \in X} u(x).$$

So a straightforward application of the results mentioned in the aside yields the following:

Corollary 2 *The preference functional U in (17) can equivalently be rewritten as follows:*

$$U(f) = (w_f(\alpha) \int_{\mathcal{X}} u(x)\, d\nu_f^\alpha(x) + (1 - w_f(\alpha)) \int_{\mathcal{X}} u(x)\, d\bar{\mu}_f^\alpha(x)), \tag{22}$$

where the integrals are taken in the Choquet sense, $w_f(\alpha)$, ν_f^α and μ_f^α are as defined above, and $\bar{\mu}_f^\alpha$ is complementary to μ_f^α .

This shows that we can always express the preferences of a DM which satisfies axioms 1-7 as the linear combination of two Choquet integrals, one with respect to a belief function, and one with respect to a capacity complementary to a belief function. The first integral reflects the pessimistic component of the decision rule, while the second reflects the optimistic component. It is immediate to verify that if $\alpha' \geq \alpha$ in a pointwise fashion then $w_f(\alpha') \geq w_f(\alpha)$. That is, the weight of the pessimistic component increases with increases of the pessimism index, as it seems natural. Two observations should be made: First, the weights of the linear combination depend on the act under consideration. Second, the belief function μ^α is in general different from ν^α. In fact it is easy to see that $\nu^\alpha = \mu^\alpha$ if and only if α is constant (see Corollary 3 below).

Remark 3 It is a simple exercise to recast the traditional "risk aversion" result in this model, as well as to verify the effects of a higher pessimism index function. Suppose that the set of consequences $\mathcal{X}$ is a subset of $\mathbf{R}$. Define, as natural, the *certainty equivalent* of act f to be the following

$$C(f) = u^{-1}(U(f)).$$

Proposition 3 *Consider two DMs, 1 and 2, with identical epistemic status $< \Pi, \mathbf{V} >$ and common beliefs given by P. Let u_i be the utility function of agent i, $i = 1, 2$, and assume it to be increasing.*[24] *Finally let α_i be the pessimism index function of i. Then the following are true:*

(i) Suppose $\alpha_1 \equiv \alpha_2$. If u_i is an increasing concave transformation of u_j (i.e., $\exists g : \mathbf{R} \to \mathbf{R}$, g concave and increasing, such that $u_i = g \circ u_j$) then for all acts $f \in \bar{\mathcal{F}}$, $C_i(f) \leq C_j(f)$;

(ii) Suppose $u_1 = u_2$. Then if $\alpha_1(X) \geq \alpha_2(X)$ for all $X \subseteq \mathcal{X}$, for all acts $f \in \bar{\mathcal{F}}$, $C_i(f) \leq C_j(f)$.

[24] I use "increasing" (resp. "decreasing") to mean non-decreasing (resp. non-increasing).

Notice that result (i) is true for all possible pessimism index functions. In fact it is easy to see that Jensen's inequality is true for Choquet integrals when integrating increasing and concave functions. Hence, for individuals with identical epistemic conditions and pessimism index, the relative concavity determines the size of the certainty equivalent of every act. Result (ii) says that, as we might expect, *ceteris paribus* a higher pessimism index function will imply lower certainty equivalents (for non-crisp acts). However the qualitative behavior of the preferences described here is different from that of a risk-averse SEU DM. That is, the preferences in this model display what after Montesano [33] and Segal and Spivak [37] is known as *first order risk aversion.* See also example 2 below. ◇

2 A Discussion

2.1 Three Special Cases

Here I will briefly discuss three cases which seem to be worthy of special attention, especially because they are formally closer to the models which have been studied in the literature.

The first obvious specialization of our representation is to the case in which the pessimism index α does not depend on X. Then we obtain a stronger version of corollary 2.

Corollary 3 *If $\succeq$ can be represented as in (16) and α is constant then U can be rewritten as follows: for all $f \in \mathcal{F}$,*

$$U(f) = \alpha \int_X u(x)\, d\nu_f(x) + (1-\alpha) \int_X u(x)\, d\bar{\nu}_f(x), \tag{23}$$

where the integrals are taken in the Choquet sense and $\bar{\nu}_f$ is complementary to ν_f.

In particular corollary 3 implies that the extremely pessimistic DM (for whom $\alpha \equiv 1$) will choose as if she evaluates the act f by a Choquet integral with respect to the supermodular distribution ν_f. On the other hand, the extremely optimistic DM ($\alpha \equiv 0$) will choose as if she evaluates each act f by a Choquet integral with respect to the submodular $\bar{\nu}_f$. The Choquet integral representation is well-known in the literature on decision theory, where it was introduced by Schmeidler [36] to model DMs who prefer to bet on events for which a probability is easily determined (like samples from an urn with known composition, as in Ellsberg's classical experiment [11]). The case where α is constant, but not necessarily 0 or 1, is instead (a generalization of) the pessimism index decision model of Arrow and Hurwicz [4].

The axiomatization of these special cases is quite straightforward, and it is especially simple in the two extreme cases. In fact then it's just a matter of reinforcing axiom 7 as follows.

Ignorance Pessimism (resp. Optimism) *Given $X, Y \in 2^{\mathfrak{X}}$, suppose that X and Y are comparable (i.e., $I(X) \cap I(Y) \neq \emptyset$) and that for every $x \in X$ there is $y \in Y$ such that $x \succeq y$ (resp. for every $y \in Y$ there is $x \in X$ such that $x \succeq y$), then $X \succeq Y$.*

In other words, a DM is extremely pessimistic (resp. optimistic) with respect to her ignorance if she evaluates sets of results by their worst (resp. best) outcome. This is clearly stronger than axiom 7. Notice that this property is required only on the ranking of "constant" acts that I defined in (3). That it holds in general for preference conditional on every non-null event descends once again from the axiom 3. Describing axiomatically a DM with constant α is not as easy. Adding the following to axiom 7 seems to provide the weakest solution.

Ignorance Substitution Consistency (ISC) *Given $X, Y \in 2^{\mathfrak{X}}$ such that $I(X) \cap I(Y) \neq \emptyset$ suppose that for some $A \subseteq I(X) \cap I(Y)$ there is $C \subseteq A$ such that if $X = \{x_1, \ldots, x_m\}$ with $x_1 \succeq \cdots \succeq x_m$, then for every $f \in \bar{\mathfrak{F}}$,*

$$(X, A;\ f, A^c) \sim (x_m, C;\ x_1, A \setminus C;\ f, A^c).$$

Then if $Y = \{y_1, \ldots, y_p\}$ with $y_1 \succeq \cdots \succeq y_p$, we have for every $f \in \bar{\mathfrak{F}}$ then

$$(Y, A;\ f, A^c) \sim (y_p, C;\ y_1, A \setminus C;\ g, A^c).$$

It is clear that Ignorance Pessimism (Optimism) also implies (in the presence of axiom 2) ISC. Notice that the latter does *not* require that the event C exists, only that if it exists, it can be used for any other set of outcomes. Also, ISC only requires such consistency to hold for some event A: As we saw earlier, axiom 3 then implies that it holds for every non-null event. I can now state the specialized version of the representation theorem.

Theorem 3 *$\succeq$ satisfies axioms 1-7 and Ignorance Substitution Consistency iff it can be represented as specified in theorem 2, where there is $\alpha \in [0, 1]$ such that for every $X \in 2^{\mathfrak{X}}$, $\alpha(X) = \alpha$. $\succeq$ satisfies axioms 1-6 and Ignorance Pessimism (resp. Optimism) iff it can be represented as specified in theorem 2, where for every $X \in 2^{\mathfrak{X}}$, $\alpha(X) = 1$ (resp. $\alpha(X) = 0$).*

It is important to observe that in all three cases the axiom is *necessary* as well as sufficient for the representation. In particular the DM whose preferences are represented by a Choquet integral with respect to ν_f must rank sets *only* on the basis of their worst outcome. One obvious conclusion to be drawn from this is that, given that the extreme pessimism associated with $\alpha \equiv 1$ seems quite excessive, we should be wary of using Choquet integrals with respect to supermodular measures as representations.[25] On the

[25] At least when unforeseen contingencies are the factor generating the non-additivity, as here. However such extreme pessimism is not so uncommon in day-to-day life. For instance, in *Uncle Vanya* Čekhov has Elena Andreyevna saying:

> After all, it seems to me that the truth, no matter what it is, is not so dreadful as uncertainty.

other hand in many circumstances it might be plausible to assume that DMs are quite (but not totally) pessimistic, i.e., they have a high α function. But, as inspection of (17) and (23) immediately confirms, the mathematical representation is continuous in α, so the behavior of the extremely pessimistic DM is after all a good approximation of the behavior of the quite pessimistic one. Dual considerations can be made for the extremely optimistic DM.

2.2 Does Ignorance Have Limits?

It is now time to present and discuss some obvious criticism to the model developed in this paper. Clearly this model suffers of all the shortcomings of the Savage model [35].[26] For instance, it is well known that axiom 2 is generally violated, especially in situations in which the assessment of probabilities of events (in Π) is based on information which the DM perceives as low quality. Clearly, a generalized model which takes care of these violations could be constructed, following the steps of Schmeidler [36] and Gilboa [18]. I decided not to do so because I wanted to focus on the effects of generalizing Savage's model to non-crisp acts, without inserting other "distorting" factors in behavior. The same goes for other violations of the Savage axioms and respective generalizations.

Let us now focus on the difference between this model and Savage's. Axiom 7 seems to me scarcely objectionable, so the axioms which call for a discussion are the ones which are crucial for the structure of the model: Axiom 3 and, to a smaller extent, axiom 4. As I mentioned earlier, theirs is the main responsibility of the state independence in the evaluation of sets of consequences. Regarding axiom 4, its role is to insure that the evaluation of the probability of event A is not affected by the fact that there is an act which yields some set X contingently on A, even if X is very large. So the probability of an event does not depend on the degree of perceived ignorance associated with that event. This seems pretty natural: In fact the degree of ignorance should in principle affect the likelihoods of the (possibly unknown to the DM) subevents of A which yield the different results but not that of A, which we assumed to be well-defined and known to the DM.

As proposition 2 makes transparent, in the set-up of this paper the state independence of axiom 3 entails a "complete ignorance" assumption: The evaluation of the (non-singleton) set of outcomes X is independent of the circumstances in which it is received, and of the act which yields it. The fact that evaluation is act-independent does not seem terribly bothersome.[27] But state independence excludes the possibility that the DM can form (possibly vague) beliefs on events she is unaware of. That is, if she can obtain $X = \{x_1, x_2\}$ contingently on event A or on event B she is not allowed to behave as if x_1

[26] Two excellent reviews are Machina [28] (for research up to 1987) and Camerer and Weber [5].

[27] Especially when taken in conjuction with axiom 2, which implies that preferences conditional on A do not depend on what could have been obtained otherwise (in the parlance of decision theorists: preferences are *consequentialist*, a term coined by Hammond [23]). For then the only important aspect of an act is its behavior conditionally on A.

is more likely than x_2 on A and x_2 is more likely than x_1 on B. The implicit assumption behind this is that if she did, she would be able to refine Π so as to write the subevents (of different conditional probability) which yield x_1 or x_2. In other words, the state space Π is the "best" state space the DM can develop, and she will not go beyond that in "rationally" assessing probabilities. This is clearly a bit extreme in drawing a drastic distinction between what is "probabilizable" and what is not.

There are however two arguments which can be offered in partial defense of axiom 3. The first is the obvious one that this model is to be understood as a first approximation, and should not be taken literally. That is, what I am modelling is a DM who has preferences which are reasonably state-independent, so that axiom 3 can be approximately taken to hold. The second more interesting point is that this way of proceeding makes the least requirements in terms of comparing theoretical (and possibly counterfactual) acts. In fact in a Savage model beliefs can only be assessed via the DM's ranking of a (large) number of theoretical acts. To illustrate this, consider a DM who *knows* that the real state space is Ω and correctly perceives every act as a function from Ω into $\mathcal{X}$. She uses instead Π, a partition of Ω, as a state space, treating acts as correspondences from Π into $\mathcal{X}$ (defined in the obvious way). Then the class of acts $\bar{\mathcal{F}}$ that the DM ranks is considerably smaller than the class of all functions from Ω into $\mathcal{X}$ that she would be asked to consider if she wanted to obey Savage's model. It is conceivable that she might find the task assigned by axiom 1 too demanding if she has to consider all possible functions on Ω, but feasible if acts are measurable with respect to Π and Π is much coarser than Ω.[28] Where does axiom 3 enter the picture? Well, suppose that X is obtained under either A or B as above. If the DM had distinct beliefs in the two cases, then she would be able to make the required comparisons, so she might as well refine the state space Π. If she *cannot* make such comparisons then clearly she does not have significantly different beliefs in the two cases, so that we can say that axiom 3 will naturally hold.[29]

For the DM studied in this paper the situation is further complicated by the fact that the DM does not even know how to refine Π. However, there is a natural way out of this: To refine each state π by considering its cartesian product with the set $\mathcal{Z} \equiv \mathcal{X}^{\mathcal{F}}$. In fact then for each $z \in \mathcal{Z}$, $f(\pi, z)$ is trivially a singleton. So if the DM can be then show to satisfy the Savage axioms on the expanded state space $\Pi \times \mathcal{Z}$ we are back into familiar terrain. But (putting aside how demanding such an expansion can be even in very simple cases) the argument I made in the last paragraph applies exactly to this proposed solution. That is, what if the DM cannot completely rank all acts on the enlarged state space? I leave it to the reader to judge whether to call this an instance of "bounded rationality". However I want to remark that the model I just outlined is *not* more general than the one presented in this paper, since it weakens axiom 3 at the cost

[28] The argument would be much more compelling if the model did not force Π to be infinite. As I mentioned in footnote 8, the extension of this model to a finite Π seems feasible and it would certainly extend its interpretational power. It is the subject of future research.

[29] Thus we see an interesting alternative interpretation of the model proposed in the paper: It is a (fairly natural) extension of Savage's model to a case in which the DM cannot make all the required rankings of theoretical acts, but can only rank a restricted set.

of strengthening axiom 1.

Having said that, it is however quite clear that it would be very interesting to obtain a more general model in which axiom 3 is significantly weakened, and α can also depend on events. This is the subject for future research.

2.3 The Related Literature

As I mentioned earlier, non-additive uncertainty is at the heart of a vast and growing literature in decision theory. This was developed in order to provide an axiomatic explanation of some commonly observed violations of Savage's Subjective Expected Utility (SEU) model, like the behavior observed in the famous Ellsberg paradox. Schmeidler's seminal work [36] proposed the non-additivity of beliefs over the state space *itself* as an explanation of the fact that subjects in Ellsberg's experiment preferred to bet on events with known odds, a phenomenon tradionally dubbed "uncertainty aversion". One intuitive justification for the non additivity is that the DM is not able to form a single-valued probability judgement for the probability of an event, but she can just assign some rough bounds for it. Subsequent works, in particular Gilboa and Schmeidler [19], Gilboa [18] and Wakker [42], have moved along this line (the latter two in a Savage framework like the one employed here). The present work moves in an orthogonal direction with respect to this literature, showing that non additivity can also be explained by the DM's perception of ignorance, what I call unforeseen contingencies. Clearly a more complete model of decision making would allow both sources of non additivity.[30]

¿From a purely formal point of view, the intuition lying behind this model is not new at all. On the one hand there is a literature dealing with complete ignorance, following Arrow and Hurwicz's seminal work [4] (e.g., Maskin [29] and Cohen and Jaffray [8]). All these works assume complete ignorance over the whole state space, that is, Π is trivial. On the other hand, ever since the first works on belief functions by Dempster [9] and Shafer [38], it has been well-known that belief functions can be justified by a "two-tiered" structure of uncertainty. One recent paper which joined these two ideas together in a fashion similar to the present one is Jaffray and Wakker [25]. Roughly stated, their idea is that the state space can be divided in two parts. One is a standard state space on which there is an "objective" probability measure. The other is a state space about which the DM is completely ignorant, in the sense that she knows its structure but is unable to specify *any* probability judgement on it. Jaffray and Wakker obtain a representation of preferences which is formally a special case of (16) but, given the differences in the set-up, has a totally different interpretation. As I discussed in footnote 21, I find the interpretation proposed here more compelling than the one they propose. Interpretation aside, there are relevant differences in the formal models: I already mentioned that they impose the result of proposition 1 as an axiom. Moreover their substitute for axiom 7

[30] Notice that the model presented here *cannot* explain the Ellsberg (and Allais) paradoxes, at least not in a straightforward manner. In fact the preferences described here satisfy the sure-thing principle, which is violated by the preferences observed in those experiments.

is a state by state dominance axiom, which would make no sense in the set-up of this model.

A similar model has been presented by Hendon, Jacobsen, Sloth and Tranæs [24], who moreover observe that an incomplete specification of the state space could explain why the DM's beliefs are represented by belief functions. They work in a generalized von Neumann-Morgenstern environment, assuming that the DM has preferences defined on a set of belief functions over $\mathcal{X}$ which obey the von Neumann-Morgenstern axioms. Comparisons between the two axiomatizations are then very hard because theirs takes advantage of the existence of an extraneous randomizing device, which I do not have. However their "consistency" axiom is very much analogous to the result of proposition 2.

The paper which is conceptually more similar to the present one is Mukerji [34], the existence of which I discovered only after formulating this model. His set-up is similar to Jaffray and Wakker's: There is a finite set of states on which the DM can formulate probability judgements, labelled Ω and a finite set of states which affect the outcomes of acts, labelled Θ. The DM has only limited capability of knowing the relation between the two sets of states, formalized by what Mukerji calls the *implication mapping*: a correspondence which associates to each $\omega \in \Omega$ a subset of Θ, which is analogous to my perception correspondence. Acts are functions from Θ to $\mathbf{R}$ (implicitly outcomes have already been converted in utilities). The epistemic construction of acts as maps from Ω to $\mathbf{R}$ is carried out as follows. Suppose that the image of ω via the implication mapping is $A \subseteq \Theta$. Then it is assumed, for an act f and a state ω, that the DM *perceives* $f(\omega)$ to be the minimum of $f(\theta)$ for $\theta \in A$. Thus an assumption of pessimism is embedded in the epistemic description of the DM, rather than described in terms of preferences as I do here. As we know from theorem 3, the representation of preferences is then a Choquet integral with respect to a belief function. So his representation is a special case of the one presented here. Moreover, even if the models are similar in their original motivation (with some slight differences in interpretation), they are technically quite different: His axiomatization follows more closely that of Schmeidler [36], and employs a primitive likelihood relation rather than deriving it from preferences for acts.

Skiadas [41] developed a very general model, which also tries to capture the problem of unforeseen contingencies. As I mentioned earlier, he departs from the standard Savage set-up in avoiding the use of consequences. He thus obtains a very general and interesting model, which can also account for some well-known violations of the Savage axioms. The model presented here has much more structure than his, and thus obtains a more precise (and restrictive) representation.[31]

Finally, the best known decision-theoretic work dealing with unforeseen contingencies is Kreps [26]. His work differs from this quite substantially. In particular by studying the DM's preferences for *opportunity sets* of acts, Kreps obtains a representation in which

[31] He observes [41, pg.28] that "providing [his utility representation] with more structure in [cases of unforeseen contingencies or when the consequences are subjective and ambiguous] is an important research challenge". The present paper moves in this direction.

the DM behaves *as if* she had her own representation of the (complete) state space in mind. Moreover she has a belief about the utility yielded by an act in each of these states, so that her perception correspondence is endogenously derived. In this respect his DM is quite different from the one modelled here: With a very rough description we could say that she copes with ignorance by creating her own world view and then behaving approximately as a SEU maximizer. In a sense this model is close to the one I discussed in the previous subsection (when defending axiom 3), and it makes similar demands (higher, actually) in terms of ranking theoretical acts. Moreover, while clearly very elegant, the endogeneity of the epistemic status of the DM turns out to be a problem in applications and extensions. How is the DM's epistemic status going to be? How will it change in dynamic problems? I feel that exogenous imposition is, in the lack of a serious cognitive model, inescapable.

3 Perception and Beliefs

3.1 The Comparative Statics of Beliefs

One of the interesting features of the model I presented in section 1 is that it can yield very simple results on the comparative statics of beliefs when the epistemic status of a DM changes. While this is an interesting question in itself, it should be observed from the outset that these results do not immediately imply very sharp predictions on behavior, as the latter often depend on the exact form and size of the DM's pessimism index $\alpha(\cdot)$. Nonetheless some results to be presented do imply predictions on behavior in the extreme cases when α is constantly equal to 0 or 1, and by continuity also in the case in which α is uniformly small or large. As these extreme cases play leading roles in the literature on applications of models with non-additive uncertainty, the exercise will hopefully turn out to be more than a pure theoretical inquiry.

It is natural to start by asking what will happen when the the DM's perception becomes richer, that is, when her perceived state space Π becomes a *finer* partition of Ω. Formally, $\Pi \uparrow \Pi'$, where each $\pi \in \Pi$ is a disjoint union of $\pi' \in \Pi'$. For instance, this is what will happen when the DM realizes that some event heretofore omitted is relevant for the decision problem, and she includes it in her description of the state space. For us this raises the important question of the way in which the DM's perception will respond to this change. That is, how is the new perception correspondence $\mathbf{V}'$ going to relate to $\mathbf{V}$, the old one? In accordance with the spirit and scope of this work I will not discuss the (extremely interesting) problem of how this adaptation is performed. Instead I will now outline some criteria that the adaptation procedure might satisfy, and show some consequences of the satisfaction of these criteria for the evolution of the DM's beliefs, i.e., her distribution functions ν_f, $f_V \in \bar{\mathfrak{F}}_V$.[32] Probably the most natural property is the

[32] Since we are going to make comparisons between DMs with different perception correspondence, I will henceforth use the $_V$ subscript when discussing acts as perceived by a DM with epistemic status $< \Pi, \mathbf{V} >$.

following.

Definition 2 *Suppose that the DM's epistemic status changes from* $< \Pi, \mathbf{V} >$ *to* $< \Pi', \mathbf{V}' >$, *where* Π *is a partition of* Π'. *We say that* $\mathbf{V}'$ *is a rationalization of* $\mathbf{V}$ *if for all* $f_V \in \bar{\mathfrak{F}}_V$ *and all* $\pi \in \Pi$,

$$(\bigcup_{\pi' \in \pi} f_{V'}(\pi')) \subseteq f_V(\pi) \tag{24}$$

with strict containment for at least one f *and one* π.

In words, the refinement in her perception of the state space allows the DM to *rationalize* some of the outcomes in $f_V(\pi)$: She understands that only a (possibly strict) subset of $f_V(\pi)$ is possible in some substate π' of π. In fact, notice that (24) is equivalent to requiring that for every $\pi' \in \pi$, $f_{V'}(\pi') \subseteq f_V(\pi)$.

In the case of a rationalization, the new information about Π can add nothing to the DM's opinion about the set of all the possible outcomes of act f. This might happen if the change in perception only satisfies the following more general property.

Definition 3 *Suppose that the DM's epistemic status changes from* $< \Pi, \mathbf{V} >$ *to* $< \Pi', \mathbf{V}' >$, *where* Π *is a partition of* Π'. *We say that* $\mathbf{V}'$ *is a revision of* $\mathbf{V}$ *if for all* $f_V \in \bar{\mathfrak{F}}_V$, *for all* $\pi \in \Pi$ *and for all* $\pi' \in \Pi$,

$$\textit{either} \quad f_{V'}(\pi') \subseteq f_V(\pi) \quad \textit{or} \quad f_{V'}(\pi') \subseteq (f_V(\pi))^c \ \textit{and} \ |f_{V'}(\pi')| = 1, \tag{25}$$

with strict containment for at least one f *and one* π'.

Clearly a rationalization is also a revision, but not conversely. The added feature of revision is that it allows the DM to "realize" that some (singleton) outcomes are altogether *absent* from her former perception $f_V(\pi)$. If we imagine the DM as facing a repeated choice problem, this is what might happen in the case of "unforeseen contingencies" *strictu sensu*: That is, when some outcomes are totally unexpected under certain circumstances. Plausibly a revision would be initiated by the occurrence of an unforeseen outcome, which then induces the DM to refine her state space in order to explain it.

Both these dynamic properties of perception have very straightforward implications for the evolution of beliefs, if we assume that the DM's likelihood judgements on the state space behave consistently. By this I mean that under the new epistemic status $< \Pi', \mathbf{V}' >$ the DM's probability distribution P' on $(\Pi', 2^{\Pi'})$ is such that for all $A \subseteq \Pi$ (which correspond to subsets in the new state space)[33]

$$P'(A) = P(A). \tag{26}$$

[33] To be rigorous I should use different symbols for the same event in the new and old state space, as strictly speaking the two sets are different objects. Since the choice of using the same symbol does not seem to entail great confusion, I opted for notational simplicity.

In other words, the added knowledge on the structure of the state space does not change the DM's beliefs on the old state space. Such assumption could be imposed behaviorally in the case of a rationalization, by additionally assuming that the ranking of crisp acts under $< \Pi, \mathbf{V} >$ (which by (24) have to be crisp also under $< \Pi', \mathbf{V}' >$) remains the same after the change in epistemic status. Since I want to focus here only on the role of modifications of the epistemic status, excluding changes in likelihood ratios (which can as well be studied in the traditional SEU framework) seems to be an obvious choice.

Proposition 4 *Suppose that* $< \Pi, \mathbf{V} > \longrightarrow < \Pi', \mathbf{V}' >$*, where* Π *is a partition of* Π'*,* $\mathbf{V}'$ *is a rationalization of* $\mathbf{V}$*, and* P' *satisfies (26). Then for all* $f_V \in \bar{\mathfrak{F}}_V$*, if* ν_f *and* ν'_f *are respectively the distribution on* $\mathfrak{X}$ *before and after the change in epistemic status,*

$$\nu'_f \geq \nu_f. \tag{27}$$

That is, $\nu'_f(X) \geq \nu_f(X)$ *for all* $X \subseteq \mathfrak{X}$*. If* Π' *is a strictly finer partition of* Ω *than* Π *and*

$$P'(\{\pi' \in \Pi' : f_{V'}(\pi') \subset f_V(\pi) \textit{ for } \pi' \in \pi\}) > 0$$

then the inequality in (27) will be strict for at least one X*.*

Proof: Notice that by (9) we have that $\nu'_f(X) = P'(A')$ where

$$A' \equiv \{\pi' \in \Pi' : f_{V'}(\pi') \subseteq X\}.$$

Analogously $\nu_f(X) = P(A)$ where

$$A \equiv \{\pi \in \Pi : f_V(\pi) \subseteq X\}.$$

As I observed earlier (24) implies that $f_{V'}(\pi') \subseteq f_V(\pi)$ for all $\pi' \in \pi$. Thus if $\pi \in A$ we immediately have that $\pi' \in A'$ for all $\pi' \in \pi$, which implies that

$$\nu'_f(X) = P'(A') \geq P(A) = \nu_f(X).$$

The last statement follows immediately from the fact that $\mathfrak{X}$ is finite. ∎

Remark 4 Another way to state proposition 4 is to say that following a rationalization the core of ν_f (as defined in (20)) "shrinks':

$$C(\nu'_f) \subseteq C(\nu_f).$$

Also, from (27) and the definition of complementary capacity (19), we immediately obtain that for all f and $X \subseteq \mathfrak{X}$,

$$\bar{\nu}'_f(X) = 1 - \nu'_f(X^c) \leq 1 - \nu_f(X^c) = \bar{\nu}_f(X). \tag{28}$$

That is, while ν_f increases (in a pointwise sense) $\bar{\nu}_f$ decreases after rationalizations. This implies that also the difference $\bar{\nu}_f - \nu_f$ decreases pointwise. ◇

When only a revision of the perception correspondence takes place we have the following result.

Proposition 5 *Suppose that* $<\Pi, \mathbf{V}> \longrightarrow <\Pi', \mathbf{V}'>$, *where* Π *is a partition of* Π', $\mathbf{V}'$ *is a revision of* $\mathbf{V}$, *and* P' *satisfies (26). Then for all* $f_V \in \bar{\mathfrak{F}}_V$, *if* ν_f *and* ν'_f *are respectively the distribution on* $\mathcal{X}$ *before and after the change in epistemic status,*

$$\nu'_f(X) + \nu'_f(X^c) \geq \nu_f(X) + \nu_f(X^c) \tag{29}$$

If moreover Π' *is strictly finer than* Π *and*

$$P'(\{\pi' \in \Pi' : f_{V'}(\pi') \subset f_V(\pi) \text{ or } f_{V'}(\pi') \subseteq (f_V(\pi))^c \text{ for } \pi' \in \pi\}) > 0$$

then the inequality will be strict for at least one X.

Proof: Let A and A' be as in the proof of proposition 4. Let

$$\tilde{A} \equiv \{\pi \in \Pi : f_V(\pi) \subseteq X^c\}$$

Define $\tilde{A}'$ analogously. For every $\pi \in A$ (25) says that for all $\pi' \in \pi$ either $\pi' \in A'$ or $\pi' \in \tilde{A}'$. Analogously for every $\pi \in \tilde{A}$ we have that every $\pi' \in \pi$ will either belong to A' or to $\tilde{A}'$. Hence,

$$\nu'_f(X) + \nu'_f(X^c) = P'(A') + P'(\tilde{A}') \geq P(A) + P(\tilde{A}) = \nu_f(X) + \nu_f(X^c).$$

As above, the last statement follows immediately from the finiteness of $\mathcal{X}$. ∎

Here we see that differently from what happens following a rationalization of perception, a revision might lead to substantial changes in the relative likelihood of a set of outcomes and its complement. On the other hand, a revision always brings forth a reduction in the non additivity of the belief function ν_f, in the following sense. Given a supermodular capacity ν defined on the subsets of a finite set S, let $N_\nu : 2^S \to \mathbf{R}$ be defined: for every $A \subseteq S$,

$$N_\nu(A) = 1 - \nu(A) - \nu(A^c). \tag{30}$$

N is a natural measure of the non additivity of a capacity at set A, and it was suggested by Dow and Werlang in [10]. It is easy to prove that if ν is a supermodular capacity then N_ν will be identically zero if and only if ν is a probability measure. Clearly (29) implies that $N_{\nu'_f} \leq N_{\nu_f}$, that is, ν'_f is everywhere more additive than ν_f. The same can be said of the complementary capacities $\bar{\nu}_f$ and $\bar{\nu}'_f$.

Given proposition 4 and remark 4 it is immediate to draw some predictions on the evolution of the preferences of pessimistic and optimistic DMs (for whom respectively $\alpha \equiv 1$ and $\alpha \equiv 0$). After a rationalization a pessimistic DM will revise upward her evaluation of non-crisp acts, while keeping fixed her evaluation of formerly crisp acts. In fact inspection of the formula of the Choquet integral (18) immediately reveals that it is increasing with upward shifts of the capacity, as in (27). Therefore her utility index for

f will increase. This is not the case for crisp acts, as their distribution does not change after rationalization. Thus after a rationalization the pessimistic DM will in a sense be more willing to take the "risk" associated with a non-crisp act (see also example 2 below). Specular considerations hold for the optimistic DM, who will see her evaluation of non-crisp acts decrease relative to crisp acts as a consequence of a rationalization. In the case of revisions, while the general reduction in non additivity discussed above does not imply specific predictions on the evolution of the DM's preferences that we have for rationalization, it can nonetheless have some interesting consequences on economic behavior, as we will see in the following simple application of these results.

Example 2 Ever since seminal work of Arrow [3, Chap. 2], it has been well-known that there is a deep connection between risk aversion and the willingness to invest in risky assets. More recently, however, it has been observed that this result relies on the particular structure of SEU preferences. For instance Arrow stated a theorem which said that a risk-neutral DM, when faced with a choice of investing either in a riskless asset, like money, or in a risky asset, will invest in the asset only if the expected return of the asset is higher than its price. Otherwise she will sell the asset short, if there is such a possibility. The relative amount of investment/sale will be determined by her risk attitude. But, as Segal and Spivak [37] and Dow and Werlang [10] observed, this is not true in models with non-additive uncertainty: These models will typically give rise, for some range of asset prices, to inertial behavior, that is, situations in which the risk-neutral DM is willing neither to buy nor to sell the stock, and she ends up investing all her wealth in money.

This is the case also for the model presented here, when we assume that the DM is pessimistic. In fact, assume that the DM has a certain wealth $W > 0$ to allocate between money, a riskless asset, and s, a risky asset. Then a result in Dow and Werlang can be restated as follows:

Proposition 6 (Dow and Werlang [10], thm. 4.2) *Suppose that the DM is pessimistic and she has an increasing and concave utility function u. She will buy a non-negative amount of the asset s if and only if its price $v \leq \int x \, d\nu_s \equiv U(s_V)$ and she will sell short a non-negative amount of it if and only if $v \geq -\int (-x) \, d\nu_s \equiv -U(-s_V)$. Such amounts will be strictly positive if the inequalities are strict.*

A more general theorem, extending to many risky assets, can be found in Simonsen and Werlang [40]. Clearly the result does not imply that there will be always portfolio inertia. In fact it is clear that if s is perceived as crisp, then ν_s will be additive and $U(s_V) = -U(-s_V)$. On the other hand the interval $[U(s_V), -U(-s_V)]$ will be non-degenerate if s is not crisp.

By using the comparative statics results we just saw and a theorem of Dow and Werlang [10] it is very easy to generalize this observation to a proposition which relates the size of the interval with the DM's perceived ignorance. The proof is immediate.

Proposition 7 *Assume that the DM is pessimistic, and that* Π' *is a refinement of* Π. *Then the following statements are equivalent:*

(i) $\mathbf{V}'$ *is a revision of* $\mathbf{V}$*;*

(ii) For all acts $f_V \in \bar{\mathfrak{F}}_V$,

$$-U'(-f_{V'}) - U'(f_{V'}) \leq -U(-f_V) - U(f_V)$$

where $U(f_V) = \int x \, d\nu_f(x)$, $-U(-f_V) = -\int (-x) \, d\nu_f(x)$ *and primed variables correspond to* $< \Pi', \mathbf{V}' >$.

The same is true for the following pair of statements:

(iii) $\mathbf{V}'$ *is a rationalization of* $\mathbf{V}$*;*

(iv) For all acts $f_V \in \bar{\mathfrak{F}}_V$,

$$[U'(f_{V'}), -U'(-f_{V'})] \subseteq [U(f_V), -U(-f_V)].$$

Thus, refinements of the state space accompanied by revisions in the perception of acts cause a reduction of the size of the sets of prices for which the DM will be inactive, though it might well be that the new interval will be quite distant from the old one. Refinements accompanied by rationalization imply that such shrinkage will occur without any significant displacement of the interval.

Notice that the considerations above are true only to the extent that the DM is quite pessimistic. In fact an optimistic DM will never be inactive: on the contrary for a range of prices she will prefer *both* buying and selling the stock to owning the riskless asset. Also it will be observed that a serious analysis of asset buying behavior should be carried out within an intertemporal model. Interestingly, in recent work Epstein and Wang [13, 14] show that models like the one presented here[34] can yield very interesting new intuitions also in the context of a more sophisticated intertemporal equilibrium model. For instance they prove in [13] that the phenomenon of portfolio inertia we just saw can give rise to an indeterminacy of the equilibrium which will actually obtain. Moreover in [14] they show that non-additivity can also explain discontinuity in equilibrium price processes, like booms and crashes in security prices, which cannot be explained by changes in fundamentals. △

A different comparative statics exercise is performed by changing the perception correspondence $\mathbf{V}$ while keeping the DM's perceived state space Π fixed. In particular one could wonder what happens when the DM realizes (possibly by personal experience) that in state π some outcomes of act f are possible which she did not acknowledge before, but she cannot explain them by changing her perceived state space Π (as she would do for a revision of her perception).

[34] They employ the model with sets of priors of Gilboa and Schmeidler [19].

Definition 4 *Suppose that the epistemic status of the DM changes from* $< \Pi, \mathbf{V} >$ *to* $< \Pi, \mathbf{V}' >$. *We say that* $\mathbf{V}'$ *is less focused than* $\mathbf{V}$ *at* $f_V \in \bar{\mathfrak{F}}_V$ *if for all* $\pi \in \Pi$

$$f_V(\pi) \subseteq f_{V'}(\pi), \tag{31}$$

with strict equality holding for at least one π.

The following result descends immediately from the definition. The proof is analogous to that of proposition 4, hence it is omitted.

Proposition 8 *Suppose that* $< \Pi, \mathbf{V} > \longrightarrow < \Pi, \mathbf{V}' >$, *where* $\mathbf{V}'$ *is less focused than* $\mathbf{V}$ *at* $f \in \bar{\mathfrak{F}}$, *and* P' *satisfies (26). Then*

$$\nu'_f \leq \nu_f, \tag{32}$$

that is, $\nu_f(X) \geq \nu'_f(X)$ *for all* $X \subseteq \mathfrak{X}$. *The last inequality will be strict for at least one* X *if moreover*

$$P(\{\pi \in \Pi : f_V(\pi) \subset f_{V'}(\pi)\}) > 0.$$

Notice that the inequality in (32) goes in the opposite direction than that in (24). That is, becoming less focused at f implies that the distribution ν_f becomes *more* non-additive and shifts downward (and specularly $\bar{\nu}_f$ shifts upward). The less focused DM, being more aware of her ignorance, has a larger "core" for her beliefs. Thus if she is pessimistic she will see her relative evaluation of act f decrease, *vice versa* if she is optimistic.

The last observation calls for an important remark on the possibilities and limits of the analysis conducted here. In general we could informally think that a DM whose perception correspondence is less focused is more "cautious", in the sense that she is more aware of her ignorance. This, however, does not imply that she will also be pessimistic. As the model clearly shows, a DM can be "cautious" and optimistic at the same time. The point is that our intuitive idea of prudence or cautiousness is really the sum of many distinct factors. There are epistemic factors which contribute to the formation of the DM's perception (which, let me stress once more, is beyond the scope of this work) and behavioral factors which bear on the DM's choice of action given her perception (which in this model are represented formally by the function α). A deeper analysis of human behavior ultimately requires the study of all these factors, and not just the behavioral ones.

3.2 On Truth and Perception

So far nothing has been said on the relation between the DM's perception and "truth", as represented by the perception of the modeller. More specifically, I implicitly assumed that the DM's perceived state space Π is a partition of the true state space Ω, so that her perception of the state space is coarse but correct. But I made no assumption on the relation between $f_V(\pi)$, the perceived set of possible results of act f in state π, and $f_U(\pi)$, the true image of f under π. The reason for this choice was that nothing in that analysis relied on the existence of a specific relationship.

On the other hand, in applications of the model to economic problems some assumptions of this type have to be made. First of all this is very important when modelling a DM facing the same decision problem repeatedly, for in such case the DM has history as a positive check on her perception. (For instance we might want to ask the obvious question: Will the DM eventually behave like a SEU maximizer?) Second, by limiting the generality of the model, these assumptions allow us to make specific refutable predictions. Clearly without making any such assumption, and in absence of "exogenous" information on the DM's epistemic status, this model could explain any observed choice pattern (even with a given Π) by an accurate choice of $\mathbf{V}$. Moreover I will argue that the conditions to be presented can be justified very naturally by simple assumptions on the dynamics of perception. Thus it is possible that a more complete model of the genesis of the epistemic "data" of the DM will eventually give them sounder foundations.

As I observed in the introduction, in the standard SEU framework we implicitly assume that the DM's is not only perfectly infomed, but also correct. That is, if x is the outcome of act f in state ω, then it is implicitly assumed that all the parties concerned will agree to that.[35] Allowing acts to be correspondences instead of functions clearly raises the question of which is the best generalization of this assumption. The following seems to be an obvious candidate.

Definition 5 *Suppose that the DM has an epistemic status given by* $< \Pi, \mathbf{V} >$. *The perception correspondence* $\mathbf{V}$ *is said to be correct at* $f_V \in \bar{\mathfrak{F}}_V$ *if for all* $\pi \in \Pi$,

$$f_V(\pi) \subseteq f_U(\pi). \tag{33}$$

$\mathbf{V}$ *is said to be accurate at* f *if (33) holds with set equality for all* $\pi \in \Pi$. *The beliefs* ν_f *induced by correct (resp. accurate) perception are called correct (resp. accurate).* $\mathbf{V}$ *is called correct (resp. accurate) if it is correct (resp. accurate) at* f *for all* $f \in \bar{\mathfrak{F}}$.

When her perception is correct, the DM might omit (or be unaware of) some outcomes, but she never believes an outcome to be possible when in fact it is not. If her perception is accurate she has a complete picture of the set of outcomes of each act in each state. Clearly such properties can be verified by the modeller who knows the DM's epistemic status, not by the DM itself.

For instance correct perception could arise if the DM faces a repeated decision problem and she forms her perception correspondence only on the basis of (reliable) historical information: Clearly if the underlying problem does not change over time, a perception thus formed must be correct. Imagine also that, unless her Π becomes finer, the DM

[35] All differences in opinion are then modelled as differences on the information about which ω really obtained.

with correct perception updates her correspondence just by adding to $f_V(\pi)$ the outcomes which she did not expect but observed (thus becoming less focused in the sense described above). Then the accuracy of her perception improves over time, and she might end up having accurate perception of some acts.

An immediate consequence of correct perception is that the set $\bar{\mathfrak{F}}_{CV}$ of (perceived) crisp acts will contain the set of Π-measurable acts, that is, acts which are (truly) constant on elements of Π. Thus different DMs with the same Π and correct, but not necessarily identical, perception correspondences agree that all Π-measurable acts are crisp. When the perception $\mathbf{V}$ is accurate at f, the distribution on $\mathfrak{X}$ induced by f_V must be equal to the following: for all $X \subseteq \mathfrak{X}$,

$$\nu_f^a(X) \equiv P(\{\pi \in \Pi : f_v(\pi) \subseteq X\}).$$

The following characterization result shows that ν_f^a, the accurate beliefs for f, play a special role. The proof follows immediately from (33) and proposition 8.

Proposition 9 *For given Π, P and $f_V \in \bar{\mathfrak{F}}_V$, suppose that $\mathbf{V}$ is correct for f. Then the belief μ on $\mathfrak{X}$ induced by f_V must satisfy*

$$\mu \geq \nu_f^a. \tag{34}$$

Vice versa the complementary capacity $\bar{\mu}$ must be pointwise lower than $\bar{\nu}_f^a$.

Remark 5 It also possible (if somewhat tedious) to prove a converse result that given any belief function μ on $\mathfrak{X}$ that satisfies (34) there is a correspondence $\mathbf{V}$ correct at f which induces μ as the distribution of f_V. This is why I called proposition 9 a characterization result. Thus proposition 9 tells us that only belief functions dominating ν_f^a can be correct beliefs for a DM with perceived state space Π. On the other hand the converse result has the less desirable implication that all such belief functions could be obtained from correct perception correspondences. ◇

Let $B_\Pi(f)$ be defined

$$B_\Pi(f) \equiv \{\mu : \mu \text{ is a belief function on } (\mathfrak{X}, 2^{\mathfrak{X}}),\ \mu \geq \nu_f^a\}.$$

Proposition 9 and remark 5 imply that $B_\Pi(f)$ is the set of correct beliefs for f under Π. If we assume that the DM's perceived state space Π becomes finer and that both the perception before the refinement, $\mathbf{V}$, and after, $\mathbf{V}'$, are accurate, then $\mathbf{V}'$ must be a rationalization of $\mathbf{V}$. Thus we have the following immediate consequence of propositions 4 and 9.

Corollary 4 *Suppose that Π becomes Π', where Π is a partition of Π', and P satisfies (26). Then the set of correct beliefs for $f \in \bar{\mathfrak{F}}$ under Π' is smaller than under Π, that is,*

$$B_{\Pi'}(f) \subseteq B_\Pi(f).$$

In the limit as $\Pi \to \Omega$ the only correct beliefs for f will be given by P_f, where for every $X \subseteq \mathfrak{X}$,

$$P_f(X) = P(\{\omega \in \Omega : f_v(\omega) \in X\}).$$

Thus we obtain an answer to a question asked at the beginning of this subsection: If the DM's beliefs are correct throughout the history of changes in her epistemic status, her beliefs will eventually be additive and accurate (even if her prior need not concide with that of the modeller). It should be stressed that unless her perception is also accurate throughout, this approach to additivity will not necessarily occur in a smooth increasing fashion. That is, she will probably undergo revisions (associated with refinements in Π) and changes in focus (with constant Π) along the way, not only rationalizations.

4 Conclusions

I have presented a model of a DM who faces a possibly coarsely specified decision problem. As the discussion in subsection 2.2 makes clear, this is by no means the unique model of such DM: It relies crucially on the "complete ignorance" assumption built into axiom 3. It is left to the reader to decide whether such assumption can be swallowed, at least as a first step. I mentioned earlier some theoretical extensions which will add generality and improve on its intuitive acceptability. On the other hand, many will agree that another important arena for testing the relevance of the model presented here is applications to some specific economic problems. Thus applications ought to have a very high priority on the research agenda opened by this work.

There is already a flourishing literature applying Schmeidler's model (or some of its relatives) to various economic phenomena, like behavior in normal-form games, the purchase of insurance or asset pricing. Some of these works have been mentioned earlier, and it is clear that *modulo* the different interpretation, they can also be seen as applications of (the pessimistic case of) this model. In a previous paper I have investigated the implications of this model and Schmeidler's for agency theory [17], especially trying to understand how the comparative statics mentioned above affect the design of incentive schemes. I am currently working on two other projects: the first is a dynamic extension of the model, with applications to questions of dynamic consistency. This exploits another feature of the structure of this model: In general the dynamic extension of models with non-additive uncertainty is not obvious (see, e.g., Gilboa and Schmeidler [20]), in the sense that there is no natural updating rule for capacities. In this model such arbitrariness seems to be mitigated, since there is a natural way to update beliefs. The second application is to the problem of incomplete contracts, which is, as I mentioned in the Introduction, certainly one of the most interesting issues that this model can address. The application would also exploit some of the specific features of the structure of the model, thus enabling us to draw some distinction between its empirical predictions and those obtained by the more traditional non-additive probability models.

One conclusion which can drawn from the work already done is that, while models of decision making with non-additive uncertainty are more complicated than the traditional SEU model, they can still be used fruitfully and with relative ease to explain many phenomena which cannot be reconciled with SEU. Hopefully the additional justification for non additivity in beliefs which this paper contributes will provide further stimulus for the study of the implications of non additivity for economic behavior.

Appendix: Proofs and Related Results

As we shall see, some of the arguments are more or less straightforward adaptation of the arguments developed by Fishburn and Savage (FS) in [16] to prove Savage's theorem in the case in which all acts are crisp.

Proof of Proposition 1: The proof makes use of the following two lemmas. The relation $\succeq_A$ is defined just after axiom 2. The proof of the first lemma is straightforward, and it is identical to the proof given by FS of the analogous result for crisp acts [16, Lemma 14.1].

Lemma 1 *Suppose that $\succeq$ satisfies axioms 1 and 2, and $\{A_1, A_2, \ldots, A_n\}$ is a partition of $A \subseteq \Pi$. If $f, g \in \bar{\mathfrak{F}}$ are such that $f \succeq_{A_i} g$ for each i then $f \succeq_A g$. If moreover $f \succ_{A_j} g$ for some j, then $f \succ_A g$.*

Lemma 2 *Suppose that $\succeq$ satisfies axioms 1-5 and $A, B \subseteq \Pi$ are such that $A \cap B = \emptyset$ and*[36] *$A \sim^* B$. Then if $f, g \in \bar{\mathfrak{F}}$ are such that for $X, Y \in 2^{\mathfrak{X}}$, $f(\pi) = X$ and $g(\pi) = Y$ for all $\pi \in A$ and $f(\pi) = Y$ and $g(\pi) = X$ for all $\pi \in B$, $f \sim_{A \cup B} g$.*

Proof: This is again similar to the result proved by FS [16, Lemma 14.2] when X, Y are both singletons. Observe that by the definition of f and g (and the fact that they are in $\bar{\mathfrak{F}}$) it must be the case that $A \cup B \subseteq C = A(X) \cap A(Y)$ and C is nonempty. Choose $x, y \in \mathfrak{X}$ such that $x \succ y$. Now, let $f', g' \in \bar{\mathfrak{F}}$ be defined as follows:

$$
\begin{aligned}
f' &= (X, A;\ Y, C \setminus A;\ y, C^c), \\
g' &= (X, B;\ Y, C \setminus B;\ y, C^c).
\end{aligned}
$$

If $X \succ Y$ (which is well-defined, since C is non-empty) then $f' \sim g'$. In fact if you let $f'' = (x, A;\ y, C \setminus A;\ y, C^c)$ and $g'' = (x, B;\ y, C \setminus B;\ y, C^c)$ then $A \sim^* B$ and (5) imply $f'' \sim g''$, and axiom 4 can be invoked to argue that then $f' \sim g'$. Clearly $f' \sim_{(A \cup B)^c} g'$ by axioms 1-2. This implies that $f' \sim_{A \cup B} g'$, for otherwise lemma 1 implies either $f' \succ g'$ or $g' \succ f'$. Since $f = f'$ and $g = g'$ on $A \cup B$, axioms 1-2 imply that $f \sim_{A \cup B} g$. The same conclusion is obtained if $Y \succ X$. Finally if $X \sim Y$ then $f \sim_{A \cup B} g$ follows immediately from axiom 3 and/or the definition of a null event and lemma 1. ∎

To prove the statement of the proposition, observe that $\nu_f = \nu_g$ is equivalent to saying that there are $\{X_1, X_2, \ldots, X_m\}$, distinct subsets of $\mathfrak{X}$, and two partitions $\{A_1, A_2, \ldots, A_m\}$ and $\{B_1, B_2, \ldots, B_m\}$ of Π such that $A_i \sim^* B_i$ and

$$
\begin{aligned}
f &= \sum_{i=1}^{m} X_i 1_{A_i} \\
g &= \sum_{i=1}^{m} X_i 1_{B_i},
\end{aligned}
$$

where 1_A is the characteristic function of event $A \subseteq \Pi$. One then shows the desired result by induction on m. Given lemmas 1 and 2, this argument is exactly like the one for crisp acts by FS [16, Theorem 14.3], so I will not repeat it here. ∎

[36] $\sim^*$ is the symmetric part of $\succeq^*$ and it is interpreted as "as likely as".

Proof of Proposition 2: Once again, the proof makes use of some simple lemmas. Before stating the first, let us agree to define an act $f \in \bar{\mathcal{F}}$ *crisp conditionally on* $A \subseteq \Pi$ if $f|_A$, the restriction of f to A, is a function. The set of acts which are crisp conditionally on A will be labelled $\bar{\mathcal{F}}_C(A)$. The set of all the probability measures on $\mathcal{X}$ is denoted $\mathcal{P}$. For any act $f \in \bar{\mathcal{F}}$ and any non-null $A \subseteq \Pi$ define the distribution induced by f *conditionally on* A to be, for every $X \subseteq \mathcal{X}$,

$$\nu_f|_A(X) = \frac{P(\{\pi \in A : f(\pi) \subseteq X\})}{P(A)}.$$

Clearly $\nu_f|_A \in \mathcal{P}$ if $f \in \bar{\mathcal{F}}_C(A)$. Finally, for $p, q \in \mathcal{P}$ and a non-null A, let

$$p \succeq_A q \Longleftrightarrow f \succeq_A g,$$

where $f, g \in \bar{\mathcal{F}}_C(A)$ are such that $\nu_f|_A = p$ and $\nu_g|_A = q$ (such acts exist by the definition of $\bar{\mathcal{F}}_C$ and the fact that P is convex-ranged). Axiom 1 and proposition 1 (which clearly holds also for restrictions to A) imply that $\succeq_A$ is well-defined and is a weak order on $\mathcal{P}$. The following lemma shows that $\succeq_A$ on $\mathcal{P}$ satisfies the independence axiom. The proof is a straightforward adaptation of FS's [16, Lemma 14.3].

Lemma 3 *Suppose that $\succeq$ satisfies axioms 1-6 and $A \subseteq \Pi$ is non-null. Then for every $p, q, r \in \mathcal{P}$ and every $\alpha \in (0,1)$*

$$p \succeq_A q \Leftrightarrow \alpha p + (1-\alpha) r \succeq_A \alpha q + (1-\alpha) r.$$

Lemma 4 *Suppose that $\succeq$ satisfies axioms 1-6 and $A \subseteq \Pi$ is non-null. If $p, q \in \mathcal{P}$ are such that $p \succ_A q$ and for $f \in \bar{\mathcal{F}}$, $g \succeq_A f \succeq_A h$ where $g, h \in \bar{\mathcal{F}}_C(A)$ are such that $\nu_g|_A = p$ and $\nu_h|_A = q$, then there is a unique $\alpha \in [0,1]$ such that*

$$f \sim_A i$$

where $i \in \bar{\mathcal{F}}_C(A)$ is such that $\nu_i|_A = \alpha p + (1-\alpha) q$.

Proof: Given proposition 1 and the observation above, we can without loss of generality abuse notation and write $\beta p + (1-\beta) q$ in place of (one of the infinitely many) acts which induce such probability distribution on $\mathcal{X}$. So we can make direct use of FS's proof of [16, Lemma 14.4] since here as there nothing is required of f except that it lies between g and h in preference. Finally, having found the α such that $\alpha p + (1-\alpha) q \sim_A f$ we can choose an act $i \in \bar{\mathcal{F}}_C(A)$ which induces such distribution conditionally on A. ∎

An immediate consequence of this result is that for a given $X \in 2^{\mathcal{X}}$ and non-null $A \subseteq I(X)$ if $x, y \in \mathcal{X}$ are such that for (all by axiom 2) $f \in \bar{\mathcal{F}}$, $(x, A;\ f, A^c) \succeq_A (X, A;\ f, A^c) \succeq_A (y, A;\ f, A^c)$ then there is $C \subseteq A$ such that

$$(X, A;\ f, A^c) \sim_A (y, C;\ x, A \setminus C;\ f, A^c). \tag{35}$$

Axiom 7 implies that for every non-null A, if $X = \{x_1, x_2, \ldots, x_m\}$ and $f \in \bar{\mathcal{F}}$ are as in the statement of the proposition, then

$$(x_1, A;\ f, A^c) \succeq_A (X, A;\ f, A^c) \succeq_A (x_m, A;\ f, A^c),$$

so that equation (35) is just equation (12). By a similar argument, for non-null $B \sim^* A$ and $g \in \bar{\mathfrak{F}}$ we can find a $D \subseteq B$ such that (13) holds. So we are only left to prove that $C \sim^* D$. For this part we need the following simple lemma, which I will also use later. But first, given the P which represents $\succeq^*$ obtained in theorem 1, let us define, for every non-null $A \subseteq \Pi$ (i.e., such that $P(A) > 0$) and every $C \subseteq \Pi$,

$$P(C|A) \equiv \frac{P(A \cap C)}{P(A)}. \tag{36}$$

Lemma 5 *Suppose that C and D are respectively subsets of non-null A and B such that (12) and (13) hold, and suppose that $\succeq$ satisfies axioms 1-6. Then if C' and D' are other subsets of A and B respectively such that (12) and (13) hold, it must be the case that $P(C') = P(C)$ (equivalently $P(C'|A) = P(C|A)$) and $P(D') = P(D)$ ($P(D'|B) = P(D|B)$).*

Proof: Given that (12) holds for both C ad C',

$$(x_m, C;\ x_1, A \setminus C;\ f, A^c) \sim (x_m, C';\ x_1, A \setminus C';\ f, A^c).$$

But if $P(C') > P(C)$ then by theorem 1, the definition of $\succeq^*$ and axiom 4 we have

$$(x_m, C;\ x_1, A \setminus C;\ f, A^c) \succ (x_m, C';\ x_1, A \setminus C';\ f, A^c),$$

yielding a contradiction. Analogously one sees that it cannot be the case that $P(C') < P(C)$. The proof that $P(D') = P(D)$ is identical modulo relabelling. ∎

Now observe that (12) is equivalent to

$$(X, A \cap B;\ X, A \setminus B;\ f, B \setminus A;\ f, (A \cup B)^c) \sim (x_m, C;\ x_1, A \setminus C;\ f, A^c), \tag{37}$$

and (13) is equivalent to

$$(X, A \cap B;\ X, B \setminus A;\ f, A \setminus B;\ f, (A \cup B)^c) \sim (x_m, D;\ x_1, B \setminus D;\ f, B^c). \tag{38}$$

But by assumption $f(\pi) = X$ for all $\pi \in (B \setminus A) \cup (A \setminus B)$. Thud the left-hand sides of the two indifferences are identical, so that axiom 1 implies

$$(x_m, C;\ x_1, A \setminus C;\ X, A \setminus B;\ f, (A \cup B)^c) \sim (x_m, D;\ x_1, B \setminus D;\ X, B \setminus A;\ f, (A \cup B)^c). \tag{39}$$

Now the fact that $A \sim^* B$ immediately implies that $(A \setminus B) \sim^* (B \setminus A)$. Suppose that $C \not\sim^* D$, say $C \succ^* D$. Find $D' \subseteq B$ such that $D' \sim^* C$ (using the convex-rangedness of P). Then proposition 1 and axiom 1 imply that D' is such that (13) holds, but $P(D') > P(D)$ a contradiction of lemma 5. The case $D \succ^* C$ cannot similarly hold, which implies that $C \sim^* D$, and concludes the proof. ∎

Proof of Corollary 1: This descends immediately from proposition 2 and the fact that $\mathcal{X}$ is finite. We can find a finite partition $\{A_1, \ldots, A_l\}$ of Π and a family $\{X_1, \ldots, X_l\}$ of subsets of $\mathcal{X}$ such that

$$f = \sum_{i=1}^{l} X_i 1_{A_i}.$$

Suppose $X_1 = \{x_1, \dots, x_p\}$, $p \geq 1$, and assume without loss of generality that $x_1 \succeq \cdots \succeq x_p$. Apply proposition 2 to A_1 and f to find $C_1 \subseteq A_1$ such that (12) holds. Let f_1 be defined as follows

$$f_1 \equiv (x_p, C_1;\ x_1, A_1 \setminus C_1;\ f, A_1^c).$$

By definition $f \sim f_1$. Construct f_2 by modifying f_1 on A_2, so that $f_1 \sim f_2$, and so on. Let $f^* \equiv f_l$, then by construction $f_* \in \bar{\mathfrak{F}}_C$ and axiom 1 implies that $f \sim f^*$. ∎

Using the definition of conditional probability in (36), we can immediately notice that in the formulation of proposition 2 we can substitute the statement $C \sim^* D$ with the equivalent $P(C|A) = P(D|B)$. As I mention in the main body of the paper, the result of that proposition can be generalized to any pair of non-null events as follows.

Theorem 4 *Suppose that $A, B \subseteq \Pi$, $f, g \in \mathcal{F}$ are such that for some $X \subseteq \mathcal{X}$, $A \subseteq f^{-1}(X)$, $B \subseteq g^{-1}(X)$ and both A and B are non-null. If $X = \{x_1, \dots, x_m\}$ relabel so that $x_1 \succeq x_2 \succeq \cdots \succeq x_m$. Then if $\succeq$ satisfies axioms 1-7 there are $C \subseteq A$ and $D \subseteq B$ such that $P(C|A) = P(D|B)$ and eqs. (12) and (13) are satisfied.*

Proof: The proof is in two steps. The first one proves the statement for pairs of events A and B with rational probabilities, the second generalizes to all pairs. The argument is not trivial, as the fact that P is only finitely additive does not permit us to use any continuity property.

Step 1 Suppose that $P(A) = r/s$ and $P(B) = p/q$. We want to prove that we can find $C \subseteq A$, $D \subseteq B$, such that (12) and (13) hold and $P(C|A) = P(D|B)$.

Using the non-atomicity of P, find finite partitions $\{A_i\}_{i=1}^{rq}$ of A and $\{B_j\}_{j=1}^{sp}$ of B such that $P(A_i) = P(B_j) = 1/sq$ for all i, j. By proposition 2 $\exists C_1 \subseteq A_1, D_1 \subseteq B_1$ such that $P(C_1) = P(D_1)$ and

$$\begin{aligned} f &\sim_A (x_m, C_1;\ x_1, A_1 \setminus C_1; f, A_1^c) &\equiv f_1, \\ g &\sim_B (x_m, D_1;\ x_1, B_1 \setminus D_1; g, B_1^c) &\equiv g_1. \end{aligned}$$

Apply proposition 2 to A_2, B_2, f_1 and g_1 to find $C_2 \subseteq A_2$ and $D_2 \subseteq B_2$ such that $P(C_2|A_2) = P(D_2|B_2)$ and

$$\begin{aligned} f_1 &\sim_A (x_m, C_2;\ x_1, A_2 \setminus C_2; f_1, A_2^c) &\equiv f_2, \\ g_1 &\sim_B (x_m, D_2;\ x_1, B_2 \setminus D_2; g_1, B_2^c) &\equiv g_2. \end{aligned}$$

Repeat this procedure recursively for $i = 1, \dots, \min\{rq, sp\}$. Say $rq < sp$, then for $sp \geq i > rq$ apply the procedure recursively letting $A_i \equiv A_{rq}$.

Proposition 2 further implies that $P(C_h|A_h) = P(C_k|A_k)$ for all $h, k \leq rq$ and $P(D_h|B_h) = P(D_k|B_k)$ for all $h, k \leq sp$. To see this, notice that lemma 5 implies that if we apply proposition 2 to A_h, A_k, f_h and f_k, and find C_h' and C_k' such that (14) and (15) hold, then $P(C_h) = P(C_h')$ and $P(C_k) = P(C_k')$. And analogously for D_h and D_k.

Let $C \equiv \cup_{i=1}^{rq} C_i$ and $D \equiv \cup_{j=1}^{sp} D_j$. Then $C \subseteq A$, $D \subseteq B$. By construction and transitivity of (conditional) indifference we have $f \sim_A f_{rq}$ and $g \sim_A g_{sp}$ which are equivalent to

$$\begin{aligned} f &\sim (x_m, C; x_1, A \setminus C; f, A^c) \\ g &\sim (x_m, D; x_1, B \setminus D; g, B^c). \end{aligned}$$

Finally

$$P(C|A) = \frac{P(\cup_i C_i)}{P(\cup_i A_i)} = \frac{\sum_i P(C_i)}{\sum_i P(A_i)} = \frac{P(C_1)}{P(A_1)} = P(C_1|A_1),$$

and similarly $P(D|B) = P(D_1|B_1)$, so that we obtain $P(C|A) = P(D|B)$. This completes the proof of step 1.

Step 2 Suppose that $P(A) = \alpha > 0$ and $P(B) = \beta > 0$, where α and β can be irrationals. Take sequences $\{\alpha_i\}$ and $\{\beta_j\}$ of rational numbers such that $\alpha_i \uparrow \alpha$ and $\beta_j \uparrow \beta$. Using non-atomicity of P, find corresponding increasing sequences of sets $\{A_i\}$ and $\{B_j\}$ such that $P(A_i) = \alpha_i$ and $P(B_j) = \beta_j$. Without loss of generality we can assume that $A_i \uparrow A$ and $B_j \uparrow B$.

Use step 1 to find $C_1 \subseteq A_1$ and $D_1 \subseteq B_1$ such that $P(C_1|A_1) = P(D_1|B_1) = k_1$ and

$$\begin{aligned} f &\sim_A (x_m, C_1;\ x_1, A_1 \setminus C_1; f, A_1^c) &\equiv f_1, \\ g &\sim_B (x_m, D_1;\ x_1, B_1 \setminus D_1; g, B_1^c) &\equiv g_1. \end{aligned}$$

Consider now $A_2' \equiv A_2 \setminus A_1$ and $B_2' \equiv B_2 \setminus B_1$. Again step 1 implies that there are C_2' and D_2' such that $P(C_2'|A_2') = P(D_2'|B_2') = k_2$ and

$$\begin{aligned} f_1 &\sim_A (x_m, C_2';\ x_1, A_2' \setminus C_2'; f, (A_2')^c) &= (x_m, C_1 \cup C_2';\ x_1, A_2 \setminus (C_1 \cup C_2');\ f, A_2^c), \\ g_1 &\sim_B (x_m, D_2';\ x_1, B_2' \setminus D_2'; g, (B_2')^c) &= (x_m, D_1 \cup D_2';\ x_1, B_2 \setminus (D_1 \cup D_2');\ f, B_2^c). \end{aligned}$$

Hence, if we let $C_2 \equiv C_1 \cup C_2'$, $D_2 \equiv D_1 \cup D_2'$, we have

$$\begin{aligned} f_1 &\sim_A (x_m, C_2;\ x_1, A_2 \setminus C_2; f_1, A_2^c) &\equiv f_2, \\ g_1 &\sim_B (x_m, D_2;\ x_1, B_2 \setminus D_2; g_1, B_2^c) &\equiv g_2, \end{aligned}$$

and

$$\begin{aligned} P(C_2|A_2) &= \frac{k_1\alpha_1 + k_2(\alpha_2 - \alpha_1)}{\alpha_2}, \\ P(D_2|B_2) &= \frac{k_1\beta_1 + k_2(\beta_2 - \beta_1)}{\beta_2}. \end{aligned}$$

We use this procedure recursively to find $C_i \subseteq A_i$ and $D_i \subseteq B_i$ such that

$$\begin{aligned} f_{i-1} &\sim_A (x_m, C_i;\ x_1, A_i \setminus C_i;\ f_{i-1}, A_i^c), \\ g_{i-1} &\sim_B (x_m, D_i;\ x_1, B_i \setminus D_i;\ g_{i-1}, B_i^c), \end{aligned}$$

and

$$\begin{aligned} P(C_i|A_i) &= \frac{\sum_{j=1}^{i} k_j(\alpha_j - \alpha_{j-1})}{\alpha_i}, \\ P(D_i|B_i) &= \frac{\sum_{j=1}^{i} k_j(\beta_j - \beta_{j-1})}{\beta_i}, \end{aligned}$$

where $\alpha_0 = 0$ and $\beta_0 = 0$. Another application of step 1 to A_i and A_{i+1}' gives as a result that $P(C_i|A_i) = P(C_{i+1}'|A_{i+1}') = k$ (the argument is analogous to the one given for step 1, so I will not repeat it here). Hence, repeating the argument for every i, we have $k_i = k$. This implies $P(C_i|A_i) = P(D_i|B_i) = k$ for all i.

Let $C \equiv \cup_{i=1}^{\infty} C_i$ and $D \equiv \cup_{i=1}^{\infty} D_i$. Clearly $C \subseteq A$ and $D \subseteq B$. Since the sequence $\{C_i\}$ is increasing to C by construction we immediately have

$$P(C_i) \le P(C) \le P(C \cup (A \setminus A_i)) = P(C_i) + (\alpha - \alpha_i).$$

Taking the limit as $i \to \infty$ on all sides we get

$$\lim_{i\to\infty} P(C_i) \le P(C) \le \lim_{i\to\infty} P(C_i) + lim_{i\to\infty}(\alpha - \alpha_i),$$

which (since the limits must exist) immediately yields

$$P(C) = \lim_{i\to\infty} P(C_i).$$

Analogously one proves that $P(D) = \lim_{i\to\infty} P(D_i)$. This immediately gives

$$P(C|A) = \frac{P(C)}{P(A)} = \frac{\lim P(C_i)}{\lim P(A_i)} = \lim \frac{P(C_i)}{P(A_i)} = \lim P(C_i|A_i) = k,$$

and analogously $P(D|B) = k$. Hence $P(C|A) = P(D|B)$.

I am only left to show that (14) and (15) hold for C and D. I prove it for C, as the proof for D is analogous. Suppose that, contradicting (14),

$$f \succ_A (x_m, C;\ x_1, A \setminus C;\ f, A^c). \tag{40}$$

Notice by construction and transitivity of indifference we have $f \sim_A f_i$ for every i. Moreover axiom 7 implies that for all i

$$f_i \preceq_A (x_m, C_i;\ x_1, A \setminus C_i;\ f, A^c),$$

so that by transitivity we obtain

$$f \preceq_A (x_m, C_i;\ x_1, A_i \setminus C_i;\ x_1, A \setminus A_i;\ f, A^c). \tag{41}$$

If C is null (implying that C_i is null for all i) this is in contradiction to eq. (40). Suppose that C is non-null. By axiom 6 given (40) we can find a finite partition $\mathcal{H}$ of A such that $\forall H \in \mathcal{H}$,

$$\begin{aligned} f \ &\succ_A \ (x_1, H;\ (x_m, C;\ x_1, A \setminus C;\ f, A^c), H^c) \\ &= \ (x_m, C \setminus H;\ x_1, A \setminus (C \setminus H);\ f, A^c). \end{aligned}$$

Given that $\mathcal{H}$ is finite we can assume without loss of generality that all $H \in \mathcal{H}$ are non-null. Since C is non-null, we can find $H \in \mathcal{H}$ such that $H \cap C$ is non-null, so that $P(C) > P(C \setminus H)$. But then there has to be i large enough so that $P(C_i) > P(C \setminus H)$, so that we immediately obtain

$$(x_m, C \setminus H;\ x_1, A \setminus (C \setminus H);\ f, A^c) \succeq_A (x_m, C_i;\ x_1, A \setminus C_i;\ f, A^c)$$

and by transitivity

$$f \succ_A (x_m, C_i;\ x_1, A \setminus C_i; f, A^c)$$

a contradiction of (41).

Suppose now that

$$f \prec_A (x_m, C;\ x_1, A \setminus C;\ f, A^c). \tag{42}$$

Again by axiom 7 we have for every i

$$f_i \succeq_A (x_m, (A \setminus A_i) \cup C_i;\ x_1, A_i \setminus C_i;\ f, A^c).$$

By transitivity

$$f \succeq_A (x_m, (A \setminus A_i) \cup C_i;\ x_1, A_i \setminus C_i;\ f, A^c). \tag{43}$$

If C is such that $P(C) = P(A)$ then it must be the case that $P(C_i) = P(A_i)$, from which we obtain an immediate contradiction of (42). Assume that $P(C) < P(A)$. Using axiom 6 once more we can find a finite partition $\mathcal{H}$ of A such that for every $H \in \mathcal{H}$,

$$\begin{aligned} f &\prec_A (x_m, H;\ (x_m, C;\ x_1, A \setminus C;\ f, A^c), H^c) \\ &= (x_m, C \cup H;\ x_1, A \setminus (C \cup H);\ f, A^c). \end{aligned}$$

We can find an H such that $P(C \cup H) > P(C)$, so that for large enough i it must be the case that $P(C \cup H) \geq P(C_i \cup (A \setminus A_i))$. This implies

$$(x_m, C \cup H;\ x_1, A \setminus (C \setminus H);\ f, A^c) \preceq_A (x_m, C_i \cup (A \setminus A_i);\ x_1, A_i \setminus C_i;\ f, A^c)$$

whence we obtain

$$f \prec_A (x_m, C_i \cup (A \setminus A_i);\ x_1, A_i \setminus C_i;\ f, A^c)$$

contradicting (43).

We can thus conclude that f satisfies (12) when C is defined as above. Repeating the argument for g and D finishes the proof of the theorem. ∎

Proof of Theorem 2: Since, as I observed earlier, Savage's axioms P1-6 hold on the set $\bar{\mathfrak{F}}_C$, it descends from Savage's representation theorem (for instance [16, Theorem 14.4]) that there is a measure P on $(\Pi, 2^\Pi)$ and a function $u : \mathcal{X} \to \mathbf{R}$ such that, for every $f, g \in \bar{\mathfrak{F}}_C$,

$$f \succeq g \iff \sum_{x \in \mathcal{X}} \varphi_f(x)\, u(x) \geq \sum_{x \in \mathcal{X}} \varphi_g(x)\, u(x). \tag{44}$$

As usual, P is unique and u is unique up to a positive affine trasformation. In lemma 1 I proved that every $f \in \bar{\mathfrak{F}}$ is indifferent to some $f^* \in \bar{\mathfrak{F}}_C$. In particular this implies

$$f \succeq g \iff f^* \succeq g^*. \tag{45}$$

Given $f \in \bar{\mathfrak{F}}$ there is a partition $\{A_1, \ldots, A_m\}$ of Π, where all A_i are non-null, and a class $\{X_i\}_{i=1}^m$ of distinct subsets of $\mathcal{X}$ such that

$$f = \sum_{i=1}^m X_i\, 1_{A_i}.$$

Let $x_i = \arg\min_{x\in X_i} u(x)$ and $x^i = \arg\max_{x\in X_i} u(x)$. We can see from the proof of corollary 1 that

$$f^* = \sum_{i=1}^{m}[x_i\, 1_{A_i\cap C_i} + x^i\, 1_{A_i\setminus C_i}]. \tag{46}$$

Applying (44) we get that the utility index for f^* is given by

$$\begin{aligned} U(f^*) &= \sum_{i=1}^{m}[P(A_i\cap C_i)\, u(x_i) + P(A_i\setminus C_i)\, u(x^i)] \\ &= \sum_{i=1}^{m} P(A_i)[P(C_i|A_i)\, u(x_i) + (1-P(C_i|A_i))\, u(x^i)] \\ &= \sum_{i=1}^{m} \varphi_f(X_i)[P(C_i|A_i) \min_{x\in X_i} u(x) + (1-P(C_i|A_i)) \max_{x\in X_i} u(x)]. \end{aligned} \tag{47}$$

Theorem 4 implies that $P(C_i|A_i)$ is a function of X_i only, so we can let $\alpha(X_i) = P(C_i|A_i)$. We can extend $\alpha(\cdot)$ to all of $2^{\mathcal{X}}$ as follows: If for $X \in 2^{\mathcal{X}}$, $I(X)$ is non-null then there is $f \in \bar{\mathfrak{F}}$ and $A \subseteq \Pi$ such that $f^{-1}(X) = A$ and A is non-null, so that $C \subseteq A$ can be found to satisfy equation (12) and $\alpha(X)$ can be defined $\alpha(X) = P(C|A)$. If $I(X)$ is null then we can let $\alpha(X) = 1$, since X cannot be the outcome of any act in $\bar{\mathfrak{F}}$ with positive probability. Thus we obtain a unique function $\alpha : 2^{\mathcal{X}} \to [0,1]$.

So we can finally rewrite (47) as follows:

$$U(f^*) = \sum_{X\subseteq\mathcal{X}} \varphi_f(X)[\alpha(X) \min_{x\in X} u(x) + (1-\alpha(X)) \max_{x\in X} u(x)],$$

which, added to (45), gives (16). Necessity of axioms 1-6 for the representation in (16) is straightforward to check. ∎

Proof of Theorem 3: The proof of the first statement is immediate given ISC. In fact axioms 1-7 imply the representation in eq. (16). Then ISC can be invoked to show that α cannot depend on the identity of X: Just consider two acts which are identical except on a set A where they respectively yield X and Y. The necessity of ISC is obvious.

As for the second statement, it is immediate to see that axiom 3 and Ignorance Pessimism imply that for every non-null $A \subseteq \Pi$ and $f \in \bar{\mathfrak{F}}$

$$(X, A;\ f, A^c) \sim_A (x_m, A;\ f, A^c).$$

Thus, in the proof of theorem 2 we obtain $\alpha(X) \equiv 1$. Necessity is obvious. As for Ignorance Optimism, the same argument proves the statement, being careful to define $\alpha(X) = 0$ for all $X \in 2^{\mathcal{X}}$ for which $I(X)$ is null. ∎

References

[1] Philippe Aghion and Benjamin Hermalin. Legal restrictions on private contracts can enhance efficiency. *Journal of Law, Economics and Organizations*, 6:381–409, 1990.

[2] Luca Anderlini and Leonardo Felli. Incompletely written contracts: Undescribable states of nature. STICERD Theor. Econ. Disc. Paper TE/93/263, London School of Economics, May 1993.

[3] Kenneth J. Arrow. *Essays in the Theory of Risk-Bearing.* North-Holland, Amsterdam, 1974.

[4] Kenneth J. Arrow and Leonid Hurwicz. An optimality criterion for decision making under ignorance. In C.F. Carter and J.L. Ford, editors, *Uncertainty and Expectations in Economics.* Basil Blackwell, Oxford, 1972.

[5] Colin Camerer and Martin Weber. Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5:325–370, 1992.

[6] Soo Hong Chew and Edi Karni. Choquet expected utility with a finite state space: Commutativity and act-independence. *Journal of Economic Theory*, 62:469–479, 1994.

[7] Gustave Choquet. Theory of capacities. *Annales de l'Institut Fourier (Grenoble)*, 5:131–295, 1953.

[8] Michèle Cohen and Jean-Yves Jaffray. Rational choice under complete ignorance. *Econometrica*, 48:1281–1299, 1983.

[9] Arthur P. Dempster. Upper and lower probabilities induced by a multi-valued mapping. *Annals of Mathematical Statistics*, 38:325–339, 1967.

[10] James Dow and S.R. da Costa Werlang. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60:197–204, 1992.

[11] Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

[12] Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, 61:1–22, 1993.

[13] Larry G. Epstein and Tan Wang. Intertemporal asset pricing under knightian uncertainty. *Econometrica*, 62:283–322, 1994.

[14] Larry G. Epstein and Tan Wang. Uncertainty, risk-neutral measures and security price booms and crashes. Working Paper 9410, Department of Economics and Institute for Policy Analysis, University of Toronto, April 1994. (to appear in *Journal of Economic Theory*).

[15] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge.* MIT Press, Cambridge, MA, 1995.

[16] Peter C. Fishburn. *Utility Theory for Decision Making.* J.Wiley and Sons, New York and London, 1970.

[17] Paolo Ghirardato. Agency theory with non-additive uncertainty. Mimeo, University of California at Berkeley, June 1994.

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

[19] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1988.

[20] Itzhak Gilboa and David Schmeidler. Updating ambiguous beliefs. *Journal of Economic Theory*, 59:33–49, 1993.

[21] Itzhak Gilboa and David Schmeidler. Additive representations of non-additive measures and the Choquet integral. *Annals of Operations Research*, 52:43–65, 1994.

[22] Faruk Gul. Savage's theorem with a finite number of states. *Journal of Economic Theory*, 57:99–110, 1992.

[23] Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.

[24] Ebbe Hendon, Hans Jørgen Jacobsen, Birgitte Sloth, and Torben Tranæs. Expected utility with lower probabilities. *Journal of Risk and Uncertainty*, 8:197–216, 1994.

[25] Jean-Yves Jaffray and Peter P. Wakker. Decision making with belief functions: Compatibility and incompatibility with the sure-thing principle. *Journal of Risk and Uncertainty*, 8:255–271, 1994.

[26] David M. Kreps. Static choice in the presence of unforeseen contingencies. In P. Dasgupta, D. Gale, O. Hart, and E. Maskin, editors, *Economic analysis of markets and games*, pages 258–281. MIT Press, Cambridge, MA, 1992.

[27] Barton L. Lipman. Limited rationality and endogenously incomplete contracts. Mimeo, Queen's University, August 1991.

[28] Mark J. Machina. Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives*, 1:121–154, 1987.

[29] Eric Maskin. Decision-making under ignorance with implications for social choice. *Theory and Decision*, 11:319–337, 1979.

[30] Salvatore Modica and Aldo Rustichini. Unawareness: A formal theory of unforeseen contingencies. Part II. Mimeo, CORE, November 1993.

[31] Salvatore Modica and Aldo Rustichini. Awareness and partitional information structures. *Theory and Decision*, 37:107–124, 1994. (formerly titled"Unawareness: A Formal Theory of Unforeseen Contingencies. Part I").

[32] Salvatore Modica, Aldo Rustichini, and Jean-Marc Tallon. A model of general equilibrium with unforeseen contingencies. Mimeo, CORE, April 1995.

[33] Aldo Montesano. The risk aversion measure without the independence axiom. *Theory and Decision*, 24:269–288, 1988.

[34] Sujoy Mukerji. Understanding the non-additive probability decision model. Mimeo, Yale University (forthcoming in *Economic Theory*), November 1993.

[35] Leonard J. Savage. *The Foundations of Statistics.* J.Wiley and Sons, New York and London, 1954.

[36] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.

[37] Uzi Segal and Avia Spivak. First order versus second order risk aversion. *Journal of Economic Theory*, 51:111–125, 1990.

[38] Glenn Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, New Jersey, 1976.

[39] Lloyd S. Shapley. Cores of convex games. *International Journal of Game Theory*, 1:11–26, 1971.

[40] Mario Henrique Simonsen and S.R. da Costa Werlang. Subadditive probabilities and portfolio inertia. *Revista de Econometria*, 11:1–19, 1991.

[41] Costis Skiadas. Conditioning and aggregation of preferences. Working Paper 142, Department of Finance, Northwestern University, November 1994. (forthcoming in *Econometrica*).

[42] Peter P. Wakker. Continuous subjective expected utility with non-additive probabilities. *Journal of Mathematical Economics*, 18:1–27, 1989.

[43] Oliver E. Williamson. *Markets and Hierarchies.* The Free Press, New York, 1975.