Incomplete preferences and confidence*

Brian Hill

HEC Paris†

July 25, 2011‡

Abstract

A theory of incomplete preferences under uncertainty is proposed, according to which a decision maker's preferences are indeterminate if and only if his confidence in the relevant beliefs does not match up to the stakes involved in the decision. We use the model of confidence in beliefs introduced in Hill (2010), and axiomatise a class of models, differing from each other in the appropriate notion of stakes. Comparative statics analysis can distinguish the decision maker's confidence from his attitude to choosing in the absence of confidence. Under one interpretation, indeterminacy of preferences can be behaviourally understood as a case where the agent defers, if a deferral option is available. We consider the case where deferral is not possible, providing an axiomatic analysis of the relationship between the decision maker's preferences in the presence and absence of a deferral option. Finally, we consider the consequences of the model in markets, where a common case of deferral is deferral to the status quo, which amounts to refusal to trade. The incorporation of confidence as proposed here appears to add extra friction, beyond the standard implications of non-expected utility models for Pareto optima.

*I would like to thank Itzhak Gilboa for valuable discussion, and the participants of D-TEA 2011 for helpful comments.

†Economics and Decision Sciences Department, HEC Paris. 1 rue de la Libération, 78351 Jouy-en-Josas, France. Tel: + 33 1 39 67 72 65. Fax: + 33 1 39 67 71 09. E-mail: hill@hec.fr.

‡This paper is a preliminary version, and may be subject to subsequent revisions. The most recent version of this paper will always be downloadable from http://www.hec.fr/hill, under the section Research, Work in progress. Please do not quote without the consent of the author.

Keywords: Incomplete preferences; confidence; multiple priors; deferral; financial markets; absence of trade.

JEL classification: D81, D01, D53.

1 Introduction

1.1 Proposal and motivation

"If you're not sure, say so". This simple maxim has mixed fortunes in decision theory. On the one hand, standard theories ignore it. They assume that, if forced to choose between two options, a decision maker will be able to choose one of them, and that this choice indicates or is determined by the decision maker's preferences over the options. Consequently, they assume that the decision maker's preferences are entirely determinate, or, as it is often put, that they are complete: he always has a (weak) preference for one of the options over the other. On the other hand, many recent theories of incomplete preferences fully embrace the maxim. In the context of decision under uncertainty, the most popular theory, the unanimity multi-prior model à la Bewley (2002), represents decision makers' beliefs by sets of probability measures, and hence can accommodate indeterminacy of probabilistic belief. Preferences are indeterminate when the beliefs do not allow a unequivocal judgement that the expected utility of one act is greater than that of the other.

However, in some ways, this theory is too naïve. First of all, according to it, the decision maker is either sure of a probabilistic belief or entirely unsure of it: it either holds for all probability measures in the set or it does not. The model cannot represent decision makers who are *more* sure of some beliefs than others, without being totally sure or unsure. Accordingly, the theory cannot allow consideration of *how* sure the decision maker has to be to form a preference. However, it seems that people often do, and should, form determinate preferences on the basis of beliefs in which there are not entirely sure in some situations – in particular, when little is at stake in the decision – whereas there are other situations – when the decision is more important, for example – in which they may need to be more sure of their beliefs to avoid indeterminacy. Indeed, the following refinement of the maxim above appears to be a more appropriate guide to when preferences are or should be indeterminate: "if you are not sure *enough*, say so".

The aim of this paper is to propose a theory of decision under uncertainty which corresponds to this refined maxim. The theory can be summed up in the following principle: one's preferences are indeterminate when and only when one's confidence in the beliefs needed to form the preference does not match up to the stakes involved in the decision. Like the standard Bewley model, it is the decision maker's beliefs which drive the indeterminacy of preferences; it is tacitly assumed that the decision maker is fully confident in his utilities. Extension to cases where confidence in beliefs and confidence in utilities are accounted for is a topic for further research.

To get a better intuition for the proposal, it is of course useful to consider behaviour that conforms to it, and this brings in the subtle question of the interpretation of incomplete preferences in choice behaviour. Although the theory proposed here applies to any interpretation, we shall take as our leading interpretation the following simple, though oft ignored way for indeterminacy of preference to come out in choices: the decision maker refrains from choosing any of the options. People sometimes have the opportunity to *defer* a decision: to abstain from choosing any of the options on offer, thus leaving the decision open, to be taken, perhaps, by their later selves, or by someone else, or by nature. In this last case, the consequence will be that which is achieved if no action is taken, and hence deferral is tantamount to choosing the status quo (when there is one). It is a natural intuition that one reason to defer is because of a lack of – or indeterminacy in – preference.[1]

Under this interpretation of indeterminacy, the theory yields an account of deferral that can be summarised in the following maxim: defer when and only when one's confidence in the beliefs needed to make a choice does not match up to the stakes involved in the choice to be made. To illustrate, consider a junior portfolio manager at a hedge fund, who currently has on his desk a portfolio consisting of stocks in a market that is not his speciality. He may be able to consult a senior, if not to just transfer the portfolio to her; this is tantamount to deferring the decision. The theory proposed here allows him to defer the decision

[1]Let us emphasise that whilst this interpretation of indeterminacy highlights one reason to defer – namely, that one is simply not sure what to choose – there are other reasons to defer – based, for example, on consideration of the utilities, information or beliefs of others, one's expectations about one's future utilities or beliefs, or one's evaluation of the status quo option – some of which have been treated elsewhere in the literature. Hence, under the interpretation in terms of deferral, the current proposal yields an account of only one aspect of deferral, which would need to be developed further to accommodate other considerations. For a discussion of the relationship to other reasons for deferral, as well as an extension to the case of deferral from menus, see Hill (2011, Section 5).

concerning a portfolio in a market with which he is not familiar whilst taking a decision on a portfolio in a market with which he is more familiar, although not an expert, but in a situation that, in all other respects, is identical to the first. In the latter situation, where he is confident enough in his judgement, it is reasonable to take the decision; however, in the former situation, where he lacks sufficient confidence, it is better to be less audacious and defer. Likewise, comparing two situations where the portfolios are in the same market with the same composition, but differ only in the importance of the client, there may be cases where the portfolio manager defers the decision concerning the prized client, but takes a decision in the case of the less important client. Whilst it would be overly cautious for him to defer when his confidence suffices for the less important decision, it would be excessively audacious to take the more important decision when this level of confidence does not match the higher stakes involved. These examples suggest that, as well as being normatively plausible, the account of incomplete preference, under the interpretation in terms of deferral, may also be descriptively reasonable in several situations.

To develop a theory of decision based on the maxim that one's preferences are indeterminate when and only when one's confidence in the beliefs needed to form the preference does not match up to the stakes involved in the decision, we use several notions introduced in Hill (2010). The decision maker's confidence in his beliefs is represented by a nested family of sets of probability measures, called a *confidence ranking*. Each set in the nested family corresponds to a level of confidence, with larger sets corresponding to higher levels of confidence. A set in the confidence ranking determines a set of probability judgements,[2] namely those which hold for every probability measure in the set. This is interpreted as the set of probability judgements that the decision maker holds with the corresponding level of confidence. Moreover, the decision maker has a function that associates to each decision a set of probability measures in the confidence ranking, called a *cautiousness coefficient*. Since a set in the confidence ranking corresponds to a level of confidence, the cautiousness coefficient can be thought of as assigning the level of confidence required in beliefs for them to play a role in the decision in question. It is defined in such a way that it picks out the same set of probability measures for all decisions having the same stakes; as such, it can be thought of as picking out the appropriate level of confidence entirely on the basis

[2]By probability judgement, we mean a statement concerning probabilities, such as "the probability of event A is greater than p".

of the stakes involved in the decision. Both of these are subjective elements in the model, representing the decision maker's confidence in his beliefs and his attitude to choosing in the absence of confidence respectively (see Section 4 below). For extended discussion of these notions, the reader is referred to Hill (2010).

As concerns stakes, one can imagine several different notions, each of which purports to determine what counts as the stakes in a decision on the basis of various aspects of that decision (see Section 2.2 for some examples). Each such notion of stakes establishes an order on decisions according to whether the stakes involved in them are higher or lower. Rather than considering a particular notion of stakes, we assume a *stakes relation* on the set of decisions that satisfies several weak properties, and perform our analysis on the basis of such a relation. Hence the representation result established here in fact applies to a family of notions of stakes, namely all those which generate stakes relations satisfying the specified properties. To the extent that each notion of stakes yields a different decision model, we thus provide an axiomatic treatment for a class of models of incomplete preferences.

In the main result of the paper, we give necessary and sufficient conditions for a representation of the following form: for all Anscombe-Aumann acts f and g, $f \preceq g$ if and only if

$$\sum_{s\in S} u(f(s)).p(s) \leqslant \sum_{s\in S} u(g(s)).p(s) \quad \text{for all } p \in D((f,g)) \tag{1}$$

where u is a von Neumann-Morgenstern utility function and D is a cautiousness coefficient for a confidence ranking Ξ. This is a representation of an incomplete preference relation $\preceq$. Under the leading interpretation suggested above, preference between elements is understood in the standard way ($f \prec g$ is interpreted as saying that the agent chooses g between f and g, and similarly for the weak order and indifference), but cases where there is no preference between acts are interpreted in terms of deferral: if neither $f \preceq g$ nor $f \succeq g$, then the decision maker defers the choice between f and g.

That this representation captures the maxim stated above can be seen as follows. First of all, note that $D((f,g))$ is the set of probability measures associated with the choice between acts f and g, and depends entirely on the stakes involved in this choice. It indicates the confidence level associated to those stakes; the only probability judgements in which the decision maker is confident enough for them to play a role in the decision are those which hold for all probability measures in $D((f,g))$. Under representation (1), g is weakly preferred to f if and only if, based only on these probability judgements, the decision maker can conclude that the expected utility of g is at least as high as that of f. So if, on the basis of these probability judgements, the decision maker can conclude neither that g has expected utility at least as high as f nor that f has expected utility at least as high as g, then he has no preference between them. In other words, his preferences over a pair of acts are indeterminate if he is not confident to the degree required by the stakes involved in the decision that one act is at least as good as the other.

Beyond relating indeterminacy of preference to confidence and the stakes involved in a choice, the model allows comparative statics analyses, which elucidate the role of the different elements in the model. In particular, from the agent's preferences at different stakes levels, one can elicit information about his confidence in his preferences. This notion can shed some light on the factors involved in an agent's decisiveness – that is, in how determinate his preferences are.

The interpretation of indeterminacy in terms of deferral leads naturally to the question of what the decision maker would do in cases where deferral is not an option. It is straightforward to extend the framework to encompass such cases, by adding a complete preference relation, representing preferences in the absence of a deferral option. By considering conditions on the relationship between preferences in the presence and absence of a deferral option, a simple representation of the latter preferences may be obtained.

Finally, the interpretation of indeterminacy of preferences in terms of deferral to the status quo amounts, in market settings, to an interpretation in terms of reluctance to trade. What consequences does representation (1) have in such settings? Consideration of this question would help put into perspective the importance of the relationship between confidence, stakes and indeterminacy of preference which underlies the representation.

The basic notions of the model are introduced and formally defined in Section 2, and the representation result is given in Section 3. In Section 4, we perform a comparative statics analysis of decisiveness. In Section 5, we consider the question of preferences in the absence of a deferral option. In Section 6, we consider the consequences of the model in markets under uncertainty. Proofs of all results are to be found in Appendix A.

In the rest of the Introduction, we discuss the relation to existing literature.

1.2 Related literature

Bewley (2002) was the first to axiomatise a "unanimity" representation of an incomplete preference relation over Anscombe-Aumann acts by a utility function and a set of probability measures, according to which there is a preference between acts if the expected utilities of the acts lie in the appropriate relation for all the probability measures in the set. To the extent that Bewley takes the strict preference relation as primitive, whereas we take the weak preference relation as primitive, our representation is technically closer to the unanimity representation used by Gilboa et al. (2010).[3] As noted above, this model cannot capture differing degrees of confidence, and hence it does not have the richness to capture the effect of the stakes involved in the decision on the degree of confidence required of beliefs to play a role in it, and on whether preference is determinate or not. The representation (1) can thus be thought of as a generalisation of the unanimity representation, replacing a single fixed set of probability measures by a family of sets, where the set of measures used varies depending on the stakes involved in the decision.

Representation (1) belongs to a family of decision models which use the representation of the agent's state of belief by a confidence ranking and are based on the idea that different sets of probability measures may be used in the evaluation of options, according to the stakes involved. This family was introduced and motivated in Hill (2010). There it was noted that members of the family differ along two dimensions: the decision rule which determines preferences on the basis of a set of probability measures and a utility function, and the notion of stakes. This paper deals mainly with the unanimity decision rule discussed above. By contrast, it remains fairly agnostic as regards the stakes: Theorem 1 axiomatises a family of models, namely those where the decision rule is the unanimity rule and the notion of stakes generates a stakes relation satisfying the basic properties laid out in Section 2.2. Moreover, Theorem 2 can be thought of as providing foundations for another class of models belonging to the same family: namely the class of models where the decision rule is the maxmin expected utility rule and the notion of stakes is as above (see Remark 3, Section 5). As such, this paper can be thought of as a complement to Hill (2010), exploring different fragments of the family introduced there.

[3]The representation in Gilboa et al. (2010) differs from representation (1) above by replacing $D((f, g))$ with a fixed set of probability measures; the representation in Bewley (2002) differs moreover in replacing the weak preferences and orders by strict preferences and orders.

Nau (1992) has proposed a theory of incomplete preferences which is similar to representation (1) in content and motivation. Besides the difference in framework (he uses the de Finetti framework), presentation (he uses confidence-weighted upper and lower conditional probabilities on random variables) and conceptualisation (the distinction between stakes and confidence is not fully brought out, and the notion of cautiousness coefficient is absent), the essential difference is that he assumes a particular notion of stakes, whereas we do not. As discussed in Section 2.2 (see in particular Remark 1), there is a sense in which his theory is a member of the family axiomatised here, and so this paper is a generalisation.

Faro (2009) has proposed an extension of Bewley's representation by incorporating a real-valued function on the space of probabilities measures, in a way inspired by the variational preferences model of Maccheroni et al. (2006). Lehrer and Teper (2011) have proposed a representation involving sets of sets of probability measures, where an act is preferred to another if it has a higher expected utility for all probability measures in at least one of the sets. There are significant differences from the current proposal in the representation of preferences (in particular, the notion of the stakes of a choice plays no role in these models), the concepts involved and how they are modelled, and the axiomatic properties (in particular, both of these models employ more severe weakenings of transitivity than used here). Faro (2009) also considers the relation to the variational preferences model, in a manner similar to our treatment of the relation to preferences in the absence of deferral in Section 5.

Seidenfeld et al. (1995); Nau (2006); Ok et al. (2008); Galaabaatar and Karni (2010) have explored extensions of Bewley's representation involving sets of probabilities and sets of utilities. As concerns the interpretation in terms of confidence in beliefs and confidence in utilities, the points made above hold in general. As stated previously, it remains to explore decision theories which incorporate confidence in beliefs and confidence in utilities.

The question of choice in the absence of deferral is technically related to Gilboa et al. (2010), Kopylov (2009) and Nehring (2009), who provide representation results on pairs of binary relations, where one is complete, the other is represented according to the unanimity representation described above, and they are represented by the same utility function (in the case of the first two papers) and related or identical sets of probability measures. Whereas Kopylov (2009) uses an interpretation of the two relations that is similar to that employed in Section 5 of this paper, Gilboa et al. (2010) and Nehring (2009) use different interpretations: they consider the relations to correspond to objective and subjective rational preferences, or beliefs and preferences respectively. Although Gilboa et al. (2010), like us, propose a representation of the second relation in terms of the maxmin expected utility representation (Gilboa and Schmeidler, 1989), Kopylov (2009) considers a special case (ϵ-contamination set of priors) and Nehring (2009) develops results for a larger class of models (invariant bi-separable preferences). All these authors work with single sets of probability measures, rather than confidence rankings. A final, axiomatic, difference is that, as well as axioms concerning the relationship between the two preference relations, they all impose axioms other than completeness on the complete preference relation (generally, at least transitivity and continuity). By contrast, here, the only axiom imposed on preferences when deferral is not an option is completeness. Accordingly, the main new axiom used here (Benchmark on certainty, Section 5) is stronger that the main "connecting" axioms used, for example, by Gilboa et al. (2010) (Caution or Default to certainty). In the case of a degenerate confidence ranking (containing a single set of probability measures), our Theorem 2 thus provides a new axiomatisation of the representation obtained in their Theorems 3 and 4.

Finally, the interpretation of indeterminacy of preference in terms of sticking to a status quo option used in Section 6 has been considered by Bewley (2002), under the name of the 'inertia assumption'. Bewley (1989) was the first to consider consequences for trade, and Rigotti and Shannon (2005) undertake a thorough analysis of markets involving decision makers with unanimity preferences. Billot et al. (2000) and Rigotti et al. (2008) consider markets involving decision makers with complete non-expected utility preferences.

2 General preliminaries

2.1 Framework

Throughout the paper, we use the standard Anscombe-Aumann framework (Anscombe and Aumann, 1963), as adapted by Fishburn (1970). Let S be a non-empty finite set of states, with Σ the algebra of all subsets of S, which are called *events*. $\Delta(\Sigma)$ is the set of probability measures on (S, Σ). Where necessary, we use the Euclidean topology on $\Delta(\Sigma)$. X is a nonempty set of outcomes; a *consequence* is a probability measure on X with finite support. $\Delta(X)$ is the set of consequences. *Acts* are functions from states to consequences; $\mathcal{A}$ is the set of acts. So, for an act f, and a state s, $f(s)$ is a lottery over X with finite support; for a utility function u over X, we will denote the expected utility of this lottery by $u(f(s)) = \sum_{x \in supp(f(s))} f(s)(x)u(x)$. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in \Re$, $0 \leqslant \alpha \leqslant 1$, the mixture $\alpha f + (1-\alpha)h$ is defined by $(\alpha f + (1-\alpha)h)(s,x) = \alpha f(s,x) + (1-\alpha)h(s,x)$ (Fishburn, 1970, Ch 13). We write $f_\alpha h$ as short for $\alpha f + (1-\alpha)h$. With slight abuse of notation, a constant act taking consequence c for every state will be denoted c and the set of constant acts will be denoted $\Delta(X)$.

We assume a preference relation on $\mathcal{A}$, denoted by $\preceq$. $\sim$ and $\prec$ are the symmetric and asymmetric components of $\preceq$, and $\asymp$ is the "determinate preference" relation, defined as follows: $f \asymp g$ iff $f \preceq g$ or $f \succeq g$.[4] So $f \not\asymp g$ means that the decision maker does not have determinate preferences between f and g; under the interpretation suggested in the Introduction, he defers the choice between these two acts.

2.2 *Stakes*

As mentioned in the Introduction, a central idea in this paper is the importance of the stakes involved in a choice for the decision made. To capture the stakes, we use a *stakes relation* that, for each pair of (binary choice) decisions the decision maker may be faced with, specifies which has higher stakes or whether the stakes involved in them are the same. That is, we assume a binary relation $\leqq$ on the set of pairs of acts ($\mathcal{A} \times \mathcal{A}$) – the *stakes relation* – which is interpreted as follows: $(f,g) \leqq (f',g')$ means that the stakes involved in the choice between f and g are (weakly) lower than the stakes involved in the choice between f' and g'. $\equiv$ and $<$ are the symmetric and asymmetric components of $\leqq$, defined in the standard way. We assume that the stakes relation satisfies the following basic properties.

(Weak Order) $\leqq$ is reflexive, transitive and complete.

(Symmetry) For all $f, g \in \mathcal{A}$, $(f,g) \equiv (g,f)$.

(Extensionality) For all $f, f', g, g' \in \mathcal{A}$, if $f(s) \sim f'(s)$ and $g(s) \sim g'(s)$ for all $s \in S$, then $(f,g) \equiv (f',g')$.

[4] Of course, if $\preceq$ is complete, then $f \asymp g$ for all f and g.

(Continuity) For all $f, f', g, g', h \in \mathcal{A}$, the sets $\{(\alpha, \beta) \in [0,1]^2 | \ (f_\alpha h, g_\beta h) \geqq (f', g')\}$ and $\{(\alpha, \beta) \in [0,1]^2 | \ (f_\alpha h, g_\beta h) \leqq (f', g')\}$ are closed in $[0,1]^2$.

(Richness) For all $f, f', g, g' \in \mathcal{A}$ such that $f(s) \not\sim g(s)$ for some $s \in S$ and $f'(s) \not\sim g'(s)$ for some $s \in S$, there exists $h, h' \in \mathcal{A}$ and $\alpha, \alpha' \in (0,1]$ such that $(f_\alpha h, g_\alpha h) \leqq (f', g') \leqq (f_{\alpha'} h', g_{\alpha'} h')$.

Weak order states that the binary choices which the agent may be faced with may be weakly ordered according to the stakes involved in them. We take this to be a basic property of the notion of stakes, and accept it without discussion here.[5]

Symmetry states that the stakes involved in a choice only depend on the alternatives available, irrespective of the order in which they are presented, and deserves no further discussion.

Extensionality states that all that counts for the stakes are the values of the consequences of the acts at the different states. If two acts are extensionally equivalent – that is, the decision maker is indifferent between the consequences of two acts at every state – then in virtually all formal theories of decision under uncertainty, they are treated (and evaluated) in exactly the same way. Extensionality says that whenever two pairs of acts are related in this way, they have the same stakes. Note that extensionality involves reference to the decision maker's preferences over constant acts. This is to be expected: the stakes involved in a choice depend on how "good" and "bad" the consequences of the various options are in particular instances, and of course, how "good" and "bad" they are depends on the decision maker's evaluation. Indeed, many plausible notions of stakes which come to mind make reference at least to the decision maker's utility function, if not also to part of his confidence ranking (see the examples below). It is thus entirely natural for the stakes relation to respect some aspects of his utility function.

On mixing a pair of acts with a third act, the stakes involved in a choice may change; Continuity says that this change is continuous in the degree of mixing. This seems reasonable: the stakes may be altered as one or both of the acts on offer are mixed with another act, but one would not expect the stakes to "jump" as the mixture coefficient moves gradually from one value to another.

[5]See Hill (2009, Section 3.2) for a discussion of the possibility of weakening the completeness property in a different but related framework.

Richness is a largely technical property, which states that the stakes involved in a choice between different acts can be shifted as far up or down the stakes order as desired, by mixing the pair of acts appropriately. There is a sense in which (in particular in the presence of an independence axiom; see Section 3) the choice between f and g and the choice between $f_\alpha h$ and $g_\alpha h$ are the "same" choice. Nevertheless, the stakes involved in these two choices may differ; to that extent, the latter choice can be thought of as a "version" of the former choice, but at the stakes level corresponding to $(f_\alpha h, g_\alpha h)$ rather than (f, g). Hence, using such mixtures, one can consider the decision maker's preferences at different stakes levels; in what follows, we will say that the decision maker prefers f to g at a certain stakes level if he prefers $f_\alpha h$ to $g_\alpha h$ for some α and h such that $(f_\alpha h, g_\alpha h)$ has that level of stakes. Richness simply states that for any (non-trivial) choice and stakes level, there is a version of the choice, obtained by mixing with a third act, which has stakes above the level in question, and there is a version which has stakes below that level. The intuition is that mixing with a third act can change many of the properties of a pair of acts, and in particular the main properties that are relevant for the stakes involved in the choice between them. Note that "trivial" choices – between an act and itself, or, given extensionality, between an act and one which is extensionally equivalent to it – have a "place" on the stakes ordering. Nevertheless, under some notions of stakes, such choices have a special status, for example as always having minimal stakes (this is the case for some of the examples mentioned below). For this reason, richness does not apply to such choices, for it unclear, on the one hand, whether the stakes in such a choice can be altered to any value by mixing and, on the other hand, whether, by mixing, one can always shift the stakes in the choice between two acts to the level of stakes involved in a choice between an act and itself.

As noted in Hill (2010, Section 4), there are many notions of stakes that can be proposed, each of which generates a particular stakes relation. The results below can be thought of as characterising a class of decision models, where the members of the class differ on the notions of stakes, but all of which use notions of stakes that generate stakes relations satisfying the five properties above. The weaker these properties, the more notions of stakes satisfy them, and the wider the results apply. To illustrate the generality of the results, and the weakness of the properties, note that the following natural notions of stakes all yield relations satisfying the properties above:[6] the stakes in the choice between f and g are given by

(i) the maximum of the negation of the utility of the least preferred consequence which could be obtained, taken over f or g

(ii) the maximum utility of the most preferred consequence which could be obtained by f or g

(iii) the maximum of the difference between the utility of the least preferred consequence and the utility of the most preferred consequence which could be obtained, taken over f or g

(iv) the probability that either of f or g takes a value below some threshold, calculated using a given probability measure (possibly taken from the decision maker's confidence ranking).

Remark 1. Consider the following notions of stakes: the stakes in the choice between f and g are given by

(v) the difference between the utility of the least preferred consequence which could be obtained by f and the utility of the least preferred consequence which could be obtained by g

(vi) the maximum absolute value of the difference between the utility of $f(s)$ and the utility of $g(s)$, taken over $s \in S$

(vii) the difference between the expected utilities of f and g, calculated using a given probability measure (possibly taken from the decision maker's confidence ranking).

It is straightforward to check that these notions generate stakes relations satisfying all of the properties above except for richness. However, they do satisfy the following weaker richness condition:

[6] Richness is satisfied by these notions of stakes under the assumption that the utility function is unbounded; the weakening of richness discussed in Remark 1 is still satisfied in the absence of this assumption.

(Richness') For all $f, f', g, g' \in \mathcal{A}$ such that $f(s) \nsim g(s)$ for some $s \in S$ and $f'(s) \nsim g'(s)$ for some $s \in S$, there exists $(f'', g''), (f''', g''') \in \|(f, g)\|$ such that $(f'', g'') \leqq (f', g') \leqq (f''', g''')$

where $\|(f, g)\|$ is the set of pairs of acts (f', g') which can either be obtained from f and g by mixing them in the same proportion with another act, or such that f and g can be obtained from f' and g' by such a mixture.[7] Rather than demanding that every (non-trivial) pair of acts can be shifted as far up or down the stakes ordering as desired by mixing with a third act, Richness' simply asks that, for every (non-trivial) pair of acts and stakes level, there is a pair of acts below (respectively, above) that level, such that either the latter pair can be obtained from the former pair by mixing, or the former can be obtained from the latter by mixing. It evidently retains the intuition behind the original richness property – in particular, the pairs in $\|(f, g)\|$ correspond to versions of the choice between f and g but with potentially different stakes. Moreover, with a corresponding modification to the axioms in Section 3, one can prove a result analogous to (but stronger than) Theorem 1 for stakes relations satisfying the first four properties above and richness'. To avoid burdening the reader with notation, we do not present such a result here, and work with the original richness property.

Note finally that the theory proposed by Nau (1992) is essentially a member of the class of representations of the form (1) which takes as stakes notion (vi). Hence the result just mentioned contains Nau's as a special case.

Remark 2. Defining the stakes relation on pairs of acts as done here is the simplest, but far from the only possibility. It could also have been defined on triples in $\mathcal{A} \times \mathcal{A} \times \Gamma$, where Γ can have several interpretations. For example, Γ can be understood as a set of context indices; hence dependence of the stakes on the context can accommodated. Alternatively, Γ could be interpreted as the status quo, if there is one; using this relation, one can capture dependence of the stakes on the status quo. It is straightforward to adapt the properties above to such stakes relations; the corresponding modifications in the axiomatisation below yield an axiomatisation of representation (1) where the stakes depend on factors other than the two acts on offer.

[7] The formal definition is as follows: for all $f', g' \in \mathcal{A}$, $(f', g') \in \|(f, g)\|$ if and only if there exists $\alpha \in (0, 1]$, $h \in \mathcal{A}$ such that $f'(s) \sim (f_\alpha h)(s)$ and $g'(s) \sim (g_\alpha h)(s)$ for all $s \in S$, or $f(s) \sim (f'_\alpha h)(s)$ and $g(s) \sim (g'_\alpha h)(s)$ for all $s \in S$.

2.3 Confidence ranking and cautiousness coefficient

Here we recall two notions that were introduced in Hill (2010).

Definition 1. A *confidence ranking* Ξ is a nested family of closed, convex subsets of $\Delta(\Sigma)$. A confidence ranking Ξ is *continuous* if, for every $C \in \Xi$, $C = \overline{\bigcup_{C' \subsetneq C} C'} = \bigcap_{C' \supsetneq C} C'$.[8] It is *centered* if $\bigcap_{\mathcal{C} \in \Xi} \mathcal{C}$ is a singleton.[9]

As mentioned in the Introduction, confidence rankings represent decision makers' beliefs, and in particular their confidence in probability judgements. The sets in the confidence ranking can be thought of as corresponding to levels of confidence. The higher the level of confidence in question, the larger the set: this translates the fact that one is confident of fewer probability judgements to that level of confidence. Confidence rankings are ordinal structures, in the sense that they correspond to specific types of weak orders on the space of probability measures (Hill, 2010, Proposition 2). The convexity and closedness of the sets of probability measures in the confidence ranking are standard assumptions in decision rules involving sets of probabilities. The continuity property guarantees a continuity in one's confidence in probability judgements: it ensures, for example, that one cannot be confident up to a certain level that probability of an event A is in $[0.3, 0.7]$ and then only confident that the probability is in $[0.1, 0.9]$ at the "next" confidence level. All confidence rankings considered in this paper will be continuous (and so continuity will often be assumed in discussion without being explicitly mentioned). Finally, a decision maker with a centred confidence ranking is one who, if forced to give his best estimate for the probability of any event, could come up with a single value (and these values satisfy the laws of probability), although he may not be very confident in it. We do not assume confidence rankings to be centred in general; in the representation result, there is an axiom that is necessary and sufficient for the confidence ranking to be centred, so this property receives a behavioural characterisation.

The second notion required is that of a *cautiousness coefficient* for a confidence ranking Ξ, which is defined to be a surjective function $D : \mathcal{A} \times \mathcal{A} \to \Xi$ that preserves $\leqq$; that is, such that for all $(f, g), (f', g') \in \mathcal{A} \times \mathcal{A}$, if $(f, g) \leqq (f', g')$, then $D((f, g)) \subseteq D((f', g'))$. D assigns to any pair of acts the level of confidence that is required in probability judgements

[8]For a set X, $\overline{X}$ is the closure of X. Note that the union of a nested family of convex sets is convex.

[9]This definition is identical to Hill (2010, Definition 2), except for a slight weakening of the notion of centeredness.

in order to be used in the choice between the acts. This level of confidence corresponds to the appropriate set of probability measures in the confidence ranking. The fact that D preserves the stakes ordering $\leqq$ implies, first of all, that D assigns a confidence level to a choice solely on the basis of the stakes involved in that choice: two pairs of acts with the same stakes (according to $\leqq$) are assigned to the same set of probability measures. Moreover, order-preservation is faithful to the intuition that the higher the stakes, the higher the confidence level required of probability judgements for them to play a role in the choice (and so the larger the relevant set of probability measures). Surjectivity of D basically attests to the behavioural nature of the confidence ranking: it implies that for each set of probability measures in the ranking, there will be a level of stakes, and hence a choice, for which it is the relevant set for that choice. For further discussion of confidence rankings, cautiousness coefficients and their properties, see Hill (2010).

3 Representation

Consider the following axioms on $\preceq$.

Axiom A1 (Determinate utilities)**.** For all $c, d \in \Delta(X)$, $c \asymp d$.

Axiom A2 (Non triviality and reflexivity)**.** $\preceq$ is non-trivial and reflexive.

Axiom A3 (Stakes-transitivity)**.** For all $f, g, h, e, e' \in \mathcal{A}$, $\alpha, \beta \in (0,1]$ such that $(f,h) \leqq (f_\alpha e, g_\alpha e)$ or $f(s) \sim g(s)$ for all $s \in S$, and $(f,h) \leqq (g_\beta e', h_\beta e')$ or $g(s) \sim h(s)$ for all $s \in S$, if $f_\alpha e \preceq g_\alpha e$ and $g_\beta e' \preceq h_\beta e'$, then $f \preceq h$.

Axiom A4 (Independence)**.** For all $f, g, h \in \mathcal{A}$ and for all $\alpha \in (0,1)$ such that $f \asymp g$ and $f_\alpha h \asymp g_\alpha h$, $f \preceq g$ if and only if $f_\alpha h \preceq g_\alpha h$.

Axiom A5 (Monotonicity)**.** For all $f, g \in \mathcal{A}$, if $f(s) \preceq g(s)$ for all $s \in S$, then $f \preceq g$.

Axiom A6 (Consistency)**.** For all $f, g, h \in \mathcal{A}$ and $\alpha \in (0,1)$ such that $(f_\alpha h, g_\alpha h) \leqq (f, g)$, if $f \asymp g$, then $f_\alpha h \asymp g_\alpha h$.

Axiom A7 (Continuity)**.** For all $f, g, h \in \mathcal{A}$, the set $\{(\alpha, \beta) \in [0,1]^2 \mid f_\alpha h \preceq g_\beta h\}$ is closed in $[0,1]^2$.

Axiom A8 (Continuity in stakes)**.** For all $f \in \mathcal{A}$ and $c, e \in \Delta(X)$ with $f \geq c > e$, if there exists $h \in \mathcal{A}$ and $\alpha \in (0,1)$ with $(f_\alpha h, e_\alpha h) > (f, c)$, then there exists $h' \in \mathcal{A}$ and $\alpha' \in (0,1)$ with $(f_{\alpha'} h', e_{\alpha'} h') > (f, c)$ such that $f_{\alpha'} h' \geq e_{\alpha'} h'$.

Axiom A9 (Centering)**.** For all $f, g \in \mathcal{A}$, there exist $h \in \mathcal{A}$ and $\alpha \in (0,1]$ such that $f_\alpha h \asymp g_\alpha h$.

Non-triviality and reflexivity (A2) and Monotonicity (A5) are entirely standard and call for no further comment. Continuity (A7) is a slight strengthening of the standard continuity axiom, and is related to axioms used elsewhere in the literature on incomplete preferences (see, for example, Dubra et al. (2004)). Indeed, in the presence of transitivity, independence and monotonicity, this axiom is equivalent to the standard one (see in particular Gilboa et al. (2010, Lemma 3)). Determinate utilities (A1), which is called C-completeness by Gilboa et al. (2010), simply says that preferences over constant acts are determinate. It translates the fact that the agent is assumed to be fully confident in his utilities; as stated in the Introduction, only confidence in beliefs is at issue here.

As concerns Stakes-transitivity (A3), note firstly that transitivity in the case of incomplete preferences involves two distinct conditions: firstly, if $f \leq g$ and $g \leq h$, then one has determinate preferences between f and h; secondly, these preferences go in the appropriate direction – that is, $f \leq h$. However, the former condition may seem too strong. For example, under the leading interpretation proposed in the Introduction, it implies that, if the decision maker makes the appropriate choices between f and g and between g and h, then he cannot defer the decision between f and h. Stakes-transitivity essentially weakens the first clause of the standard transitivity property, whilst retaining its second clause. More precisely, except for cases where f and g or g and h are extensionally equivalent, it demands determinate preference between f and h only when the decision maker's preferences between f and g and between g and h are determinate for stakes higher than the stakes in the choice between f and h;[10] it allows preferences to be indeterminate if this is not the case. In terms of the deferral interpretation, stakes-transitivity allows the decision maker to defer the choice between f and h even if he chooses f over g and g over h; but this is allowed only when one of the latter two decisions is less important than the former one. A decision maker may prefer spending \$10 on a bet on a certain ambiguous event to

[10] Recall from the discussion of richness in Section 2.2 that preferences at different stakes levels are given by preferences over appropriate mixtures of the acts in question.

his current portfolio, no matter what his current portfolio is; transitivity (applied repeated) implies that he prefers spending $10 000 on 1000 bets on this same event to his current portfolio, whereas stakes-transitivity allows indeterminacy of preferences for this choice, insofar as the stakes are higher than for a single $10 bet. Moreover, the axiom implies that, whenever preferences are determinate, they go in the direction implied by the standard transitivity axiom. So, to the extent that one can speak of "violations" of the standard transitivity axiom, they never result in acyclic preferences, but only in indeterminacy of preference between options where transitivity would have implied a determinate preference. Indeed, if preferences are complete, stakes-transitivity is equivalent to transitivity, in the presence of independence.[11]

A similar situation holds for Independence (A4). Whereas the standard independence axiom implies, firstly, that certain preferences are determinate, and secondly, they go in a certain direction, the independence axiom used here simply states that whenever preferences are determinate, they go in the direction specified by the standard independence condition. Evidently it fully retains the intuitions of the standard axiom, and is equivalent to it in the case of complete preferences; indeed, it can be thought of as an alternative way of extending the traditional independence axiom to the case of incomplete preferences, which separates the part of the standard axiom concerning determinacy of preference from the part concerning direction of preference.

Consistency (A6) is perhaps the most novel axiom and naturally so: it deals with the relation between choices at different stakes levels. It says that, if preferences between f and g are determinate, then preferences will be determinate between any mixtures of f and of g, as long as the stakes are not higher. In other words, if one has determinate preferences between two options at a given stakes level, then as the stakes fall, one retains the determinacy of the preferences. If you can choose between the options when there are hundreds of thousands of dollars at stake, then one can still choose when there are only

[11] As anticipated in Section 1.2, the mildness of this weakening of transitivity is a major difference between the current model and other models of incomplete preferences that seek to go beyond the unanimity model à la Bewley. For example, one can imagine a model with a second-order probability measure instead of a confidence ranking, and according to which there is a preference between acts if the probability that the expected utility of one act is greater than that of the other is above a certain threshold; such a model involves more serious violations of the transitivity axiom. Riella and Teper (2011, Theorem 2) axiomatise a model similar to that just described, with the threshold applied to first-order rather than second-order probabilities.

tens of thousands at stake. As such, it is this axiom in particular which translates the idea that the higher the stakes, the more confidence is needed to take the choice. This is a fully behavioural axiom, which is in principle testable by, for example, comparing decisions to defer at different stakes levels. (Of course, the other axioms discussed so far are as behavioural as their standard counterparts.)

Of the final two axioms, Continuity in stakes (A8) is a largely technical axiom. It states that, whenever an act f is preferred to a constant act c, then as the stakes in the choice are gradually increased (supposing they are not maximal), the act may no longer be preferred to c, but the most preferred constant act to which it is preferred will not suddenly "jump" down with a slight increase in stakes. In other words, this axiom states that the set of constant acts in the lower contour set of an act changes gradually with an increase in the stakes involved in the decision.

The final axiom, Centering (A9), is not required for the representation in general, but corresponds to the centeredness property of the confidence ranking. It states that, for any pair of acts, if the stakes are low enough in the choice between them, the decision maker's preferences over them are determinate. If there is sufficiently little at stake, the decision maker will hazard a choice, in which he may not be very confident.

These axioms are necessary and sufficient for the desired representation of preferences.

Theorem 1. *The following are equivalent.*

(i) $\preceq$ *satisfies A1-A8,*

(ii) there exists a nonconstant utility function $u : X \to \Re$*, a continuous confidence ranking* Ξ *and a cautiousness coefficient* $D : \mathcal{A} \times \mathcal{A} \to \Xi$ *such that, for all* $f, g \in \mathcal{A}$*,* $f \preceq g$ *iff*

$$(1) \qquad \sum_{s \in S} u(f(s)).p(s) \leqslant \sum_{s \in S} u(g(s)).p(s) \quad \textit{for all } p \in D((f,g))$$

Furthermore, for any other nonconstant utility function u'*, continuous confidence ranking* Ξ' *and cautiousness coefficient* D' *representing* $\preceq$ *according to (1), there exists a positive real number* a *and a real number* b *such that* $u' = au + b$*,* $\Xi' = \Xi$ *and* $D' = D$*.*

Finally, under the conditions above, A9 is satisfied if and only if Ξ *is centred.*

4 Attitudes to choosing in the absence of confidence

In this section, we undertake a basic comparative statics analysis for the model proposed above. Throughout the section, we assume a fixed stakes relation $\leqq$.

Models of incomplete preferences admit a particularly simple comparison of decision makers, in terms of the completeness of the preference relation. For decision makers 1 and 2, if 2 (weakly) prefers f to g whenever 1 does, this is an indication that 2 is less prone to indeterminacy of preference than 1. In terms of the leading choice-based interpretation of indeterminacy proposed in the Introduction, this indicates that 1 is more prone to defer, and less prone to decide, than 2. Formally, we say that $\preceq^1$ is *less decisive* than $\preceq^2$ if $\preceq^1 \subseteq \preceq^2$.[12]

In order to characterise this relation in terms of the elements of the model, we require the relation $\sqsubseteq$ on families of sets, defined as follows. For two families of sets Ξ and Ξ', we write $\Xi \sqsubseteq \Xi'$ when, for every $\mathcal{C} \in \Xi$, there exists $\mathcal{C}' \in \Xi'$ with $\mathcal{C} \subseteq \mathcal{C}'$. We have the following result.

Proposition 1. *Suppose that $\preceq^1$ and $\preceq^2$ satisfy axioms A1–A8, and are represented by (u_1, Ξ_1, D_1) and (u_2, Ξ_2, D_2) respectively. The following are equivalent:*

(i) $\preceq^1$ *is less decisive than* $\preceq^2$

(ii) u_2 *is a positive affine transformation of* u_1*,* $\Xi_2 \sqsubseteq \Xi_1$ *and* $D_2((f,g)) \subseteq D_1((f,g))$ *for all* $(f,g) \in \mathcal{A}^2$.

Besides the utility functions, the main two elements in this result are the confidence ranking and the cautiousness coefficient. However, they may be understood as playing different roles. Consider two decision makers with the same utility function. Decision maker 1 has unanimity preferences à la Bewley: his confidence ranking contains a single set of probability measures $\mathcal{C}$ and the cautiousness coefficient sends all pairs of acts to that set. Decision maker 2, by contrast, has a rich confidence ranking Ξ, with many sets of probability measures, and an appropriate cautiousness coefficient, sending different pairs of acts to different sets. As long as $\mathcal{C}' \subseteq \mathcal{C}$ for all $\mathcal{C}' \in \Xi$, 1 will be less decisive than 2. However, there seems to be something more precise to say about the relationship between the two decision makers. In particular, it appears that, on the one hand, 1 is less sensitive

[12] Containment of the preference relations is equivalent to saying that, for all $f, g \in \mathcal{A}$, if $f \succeq^1 g$ then $f \succeq^2 g$.

to the importance of decisions than 2 – if he prefers f over g at a given stakes level, then he has the same preference at any stakes level, no matter how high – but, on the other hand, 1 is confident of fewer beliefs than 2. In other words, there seems to be an aspect of beliefs (how confident one is of certain beliefs) as well as an aspect of taste (how willing one is to decide on the basis of beliefs in which one has a certain amount of confidence) mixed together in Proposition 1. To tease them apart, let us introduce the notion of confidence in preferences.

Definition 2. Let $\preceq$ satisfy axioms A1–A8. We define the *confidence in preferences* relation $\preccurlyeq$ on $\mathcal{A}^2$ as follows: for any $f, g, f', g' \in \mathcal{A}$, $(f, g) \preccurlyeq (f', g')$ iff, for all $\alpha, \alpha' \in (0, 1], h, h' \in \mathcal{A}$ such that $(f_\alpha h, g_\alpha h) \equiv (f'_{\alpha'} h', g'_{\alpha'} h')$:

$$f_\alpha h \succeq g_\alpha h \Rightarrow f'_{\alpha'} h' \succeq g'_{\alpha'} h'.$$

Definition 2 relies on the observation that, if a decision maker prefers f' to g' at a given stakes level but has indeterminate preferences between f and g at that level, then this can be taken as an indication that he is *more confident in his preference for f' over g' than in his preference for f over g.*[13] In other words, one can extract information about a decision maker's confidence in his preferences from the extent to which he hold specific preferences at given levels of stakes. This is done according to the simple principle: the preferences that the decision maker holds at higher stakes are those in which he is more confident.

Given these considerations, we shall say that two decision makers are confidence equivalent if they have the same confidence in preferences.

Definition 3. Let $\preceq^1$ and $\preceq^2$ be preference relations satisfying the axioms A1–A8. $\preceq^1$ and $\preceq^2$ are *confidence equivalent* if $\preccurlyeq^1 = \preccurlyeq^2$.

Proposition 2. *Let $\preceq^1$ and $\preceq^2$ be preference relations satisfying the axioms A1–A8, and represented by utility functions, confidence rankings and cautiousness coefficients (u_1, Ξ_1, D_1) and (u_2, Ξ_2, D_2) respectively. $\preceq^1$ and $\preceq^2$ are confidence equivalent iff u_2 is a positive affine transformation of u_1, and $\Xi_1 = \Xi_2$.*

[13] Recall from the discussion of richness in Section 2.2 that talk of preferences at different stakes levels is spelt out formally in terms of preferences over appropriate mixtures.

This proposition confirms that a decision maker's confidence in his preferences is entirely determined by his utilities over consequences and his belief state (including his confidence in beliefs). This is to be expected: to the extent that preferences are determined by utilities and beliefs, it is reasonable that confidence in preferences be determined by confidence in utilities – which is trivial in this model, because of the use of a single utility function – and confidence in beliefs, represented by the confidence ranking. The notion of confidence in preferences also helps shed light on the example above: decision makers 1 and 2 obviously have different confidence in preferences, and it is this difference, as much as any difference in attitude to choosing in the absence of confidence, that yields the difference in decisiveness. In fact, once differences in confidence in beliefs are accounted for, by comparing decision makers who have the same confidence in preferences, decisiveness is entirely determined by the cautiousness coefficient, as the following corollary of Propositions 1 and 2 shows.

Corollary 1. *Suppose that* $\preceq^1$ *and* $\preceq^2$ *satisfy axioms A1–A8, are confidence equivalent, and are represented by* (u, Ξ, D_1) *and* (u, Ξ, D_2) *respectively. The following are equivalent:*

(i) $\preceq^1$ *is less decisive than* $\preceq^2$

(ii) $D_2((f, g)) \subseteq D_1((f, g))$ *for all* $(f, g) \in \mathcal{A}^2$.

In summary, one can compare decision makers' decisiveness, that is, the pairs of acts for which they have determinate preferences. In the absence of any condition on the decision makers, a decisiveness ordering corresponds to identical utilities (up to positive affine transformation) and appropriate orders on the confidence rankings and cautiousness coefficients. Furthermore, one can derive a decision maker's confidence in preferences from his preferences. Decision makers with the same confidence in preferences are precisely those who share the same utilities (up to positive affine transformation) and confidence in beliefs, represented by a confidence ranking. Finally, for decision makers with the same confidence in preferences, comparison in terms of decisiveness corresponds precisely to the appropriate relation between their cautiousness coefficients. To the extent that such decision makers have the same confidence, differences in decisiveness between them must come down to differences in their attitudes to choosing on the basis of limited confidence. Accordingly, a decision maker's cautiousness coefficient captures precisely his attitude to choosing in the absence of confidence.

5 Absence and presence of the option to defer

In the previous sections, we have considered incomplete preferences, and proposed a choice-theoretic interpretation of indeterminacy of preference as taking an option to defer – to someone else, to one's later self, or to a status quo – that was assumed to be always present. In reality, of course, deferral is not always an option: it is not always true that one can leave the decision to someone else or until later, or that there is always a status quo. It is thus natural to ask, given the decision maker's preferences in the presence of a deferral option, what can be said about his preferences when deferral is not an option. In this section, we propose a framework for considering this question, and provide a simple axiomatic answer.

In order to consider choices in the presence and absence of the option to defer, the most natural framework consists of two binary relations, $\preceq^d$ and $\preceq^n$, over the set of acts. The relation $\preceq^d$ represents the decision maker's preferences in the presence of a deferral option, whereas the relation $\preceq^n$ represents his preferences in the absence of a deferral option. The issue of his preferences in the presence and absence of deferral, as well as of the relationship between these preferences, can be tackled by considering axioms on the preference relations and the relationship between them, such as the following.

Axiom B1 (Deferral)**.** $\preceq^d$ satisfies A1-A8.

Axiom B2 (Forced choice)**.** $\preceq^n$ is complete.

Axiom B3 (Benchmark on certainty)**.** For all $f, g \in \mathcal{A}$, if there is no $c \in \Delta(X)$ such that $f_\alpha h \succeq^d c_\alpha h$ but $g_{\alpha'} h' \npreceq^d c_{\alpha'} h'$ for some $h, h' \in \mathcal{A}$ and $\alpha, \alpha' \in (0, 1]$ with $(f_\alpha h, c_\alpha h) \equiv (g_{\alpha'} h', c_{\alpha'} h') \equiv (f, g)$, then $g \nprec^n f$.

The first two axioms state basic properties of the preferences in the presence and absence of a deferral option. Deferral (B1) states that the decision maker's preferences in the presence of a deferral option satisfy the axioms proposed and discussed in Section 3. Forced choice (B2) states that, in the absence of deferral, preferences are complete. Under the leading interpretation proposed above, cases of indeterminacy of preference are understood as cases where the decision maker defers; by demanding that the preferences in the absence of deferral ($\preceq^n$) are complete, this axiom translates the fact that deferral is not possible.

The final axiom, Benchmark on certainty (B3), concerns the relationship between the decision maker's preferences in the presence and absence of a deferral option. Preferences in the presence of a deferral option provide a crude indication of the relative worth of different acts (for the decision maker); one way of getting a more refined judgement is by considering how the acts compare to constant acts. Thus, even if the decision maker would defer the choice between two acts, say f and h, he may decide in favour of h over a particular constant act (for example, a sure \$5), whilst not choosing f over that act (in cases where these decisions have the same stakes as the decision between f and h). Benchmark on certainty demands that only in such cases may he strictly prefer h over f when deferral is not an option. In other words, it states that if an act does not fair favourably with respect to another when compared to constant acts and deferral is an option, it cannot be strictly preferred when deferral is not an option. This axiom characterises a form of conservative decision making, insofar as only the constant acts that the decision maker considers worse than the acts in question when deferral is an option count in determining preferences when deferral is not an option.

These axioms are necessary and sufficient for the following joint representation of preferences in the absence and presence of deferral.

Theorem 2. *The following are equivalent.*

(i) $\preceq^d$ *and* $\preceq^n$ *satisfy B1-B3,*

(ii) there exists a nonconstant utility function $u : X \to \Re$*, a continuous confidence ranking* Ξ *and a cautiousness coefficient* $D : \mathcal{A} \times \mathcal{A} \to \Xi$ *such that, for all* $f, g \in \mathcal{A}$*,*

(a) $f \preceq^d g$ *iff*

$$\sum_{s \in S} u(f(s)).p(s) \leqslant \sum_{s \in S} u(g(s)).p(s) \quad \text{for all } p \in D((f,g)) \tag{1}$$

(b) $f \preceq^n g$ *iff*

$$\min_{p \in D((f,g))} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in D((f,g))} \sum_{s \in S} u(g(s)).p(s) \tag{2}$$

Furthermore, for any other nonconstant utility function u'*, continuous confidence ranking* Ξ' *and cautiousness coefficient* D' *representing* $\preceq^d$ *and* $\preceq^n$ *according to (1) and (2), there exists a positive real number* a *and a real number* b *such that* $u' = au + b$*,* $\Xi' = \Xi$ *and* $D' = D$*.*

The axioms imply a simple relationship between the decision maker's preferences in the presence and absence of a deferral option. Whereas, in the presence of a deferral option, the decision maker prefers an act to another if and only if it has higher expected utility according to all the probability measures in the set corresponding to the level of confidence appropriate for the stakes involved in the choice, in the absence of deferral, it is the relationship between the minimum expected utilities over the same set of probability measures which determines his preference. In other words, whereas in the presence of deferral, he prefers an act to another if he is sufficiently confident in the relevant probability judgements, when he is forced to decide, he chooses the act that has the highest worst-case expected utility, given the judgements in which he has sufficient confidence.

As discussed in Section 1.2, the representation, as well as the setup in terms of two preference relations, is similar to that proposed by Gilboa et al. (2010). As they note, replacing the main axiom – in the current case, Benchmark to certainty (B3) – by other axioms may yield representations using sets of probability measures which differ from those used in the evaluation of preferences in the presence of deferral. Similarly, one may obtain representations involving the same set of probability measures, but different rules, such as the maxmax expected utility rule or the α-maxmin expected utility rule. To this extent, Theorem 2 can be thought of as one of a family of representations of the relationship between preferences in the presence and in the absence of deferral; exploration of other members of the family is a task for future research.

Remark 3. As Gilboa et al. (2010) note, their result can be thought of as providing a novel foundation for the maxmin representation of Gilboa and Schmeidler (1989). Similarly, Theorem 2 provides foundations for a class of preferences which belong to the family identified in Hill (2010), namely those preferences using the maxmin expected utility decision rule and a notion of stakes which yields a stakes relation satisfying the properties given in Section 2.2. This is the first axiomatisation of this class of preferences, to the knowledge of the author.[14]

[14] Note in particular that the model considered in Hill (2010), although it uses the same decision rule, does not belong this class: there the stakes are a function of a single act, whereas here the stakes are a function of pairs of acts.

6 Confidence and indeterminacy in markets

As a further exploration of the wider implications of the model, we briefly consider some consequences for risk sharing in financial markets. Recall the leading interpretation of indeterminacy of preferences in terms of deferral, where one possible sort of deferral consists of taking a status quo option, if present. A status quo is present in a market setting: it is simply the option of not trading. Bewley (1989) and Rigotti and Shannon (2005) have considered the consequences of the unanimity model (à la Bewley) in a market setting, interpreting indeterminacy of preferences as the choice of not trading; we do the same here for preferences represented according to (1).

We consider a standard Arrow-Debreu exchange economy with a complete set of (non-negative) state-contingent commodities on a finite state space S. The set of acts $\mathcal{A}$ is defined as in Section 2.1, with the set of outcomes specified by $X = \Re_+$. A state-contingent commodity is a vector in $\Re_+^S$, and can be naturally assimilated with the corresponding element in $\mathcal{A}$.[15] The economy has finitely many agents, indexed by $i = 1 \ldots n$. Each has preferences $\preceq^i$ over $\mathcal{A}$ (and hence over $\Re_+^S$), represented as in (1) for a stakes relation $\leqq^i$. Each agent has thus a utility function $u^i : \Re_+ \to \Re$, a confidence ranking Ξ^i on S and a cautiousness coefficient D^i. We assume that all u^i are differentiable, strictly concave and strictly increasing. Note that, since expected utility preferences and unanimity preferences à la Bewley are special cases of (1), the economy may contain agents with these sorts of preferences. The aggregate endowment is $e \in \Re_{++}^S$. Finally, an allocation $(x^1, \ldots, x^n) \in (\Re_+^S)^n$ is said to be *feasible* when $\sum_i x^i = e$, and it is *interior* if $x_s^i > 0$ for all i and s.

Definition 4. An allocation $(y^1, \ldots, y^n)$ *Pareto dominates* the allocation $(x^1, \ldots, x^n)$ if, for each agent i, either $y^i \succ^i x^i$ or $y^i = x^i$.

A feasible allocation $(x^1, \ldots, x^n)$ is *Pareto optimal* if there is no feasible allocation that Pareto dominates it.

[15] State-contingent commodities correspond to acts whose consequences are degenerate lotteries; $\Re_+^S$ thus corresponds to a proper subset of $\mathcal{A}$. Nevertheless, thanks to the continuity of the utility function, every lottery over $\Re_+$ (ie. element of $\Delta(X)$) has a certainty equivalent in $\Re_+$, so, given the utility function, preferences over $\mathcal{A}$ are completely determined by preferences over $\Re_+^S$. Hence, although we assume preferences over $\mathcal{A}$, this is equivalent in this setup to assuming preferences over $\Re_+^S$; similarly, properties of preferences can be formulated either in terms of $\mathcal{A}$ or $\Re_+^S$. We continue to use the notation introduced in Section 2, and in particular the generic symbols $f, g \ldots$ for acts; we use standard vector notation and generic symbols $x, z \ldots$ for commodities.

This notion of Pareto optimality is very close to that studied by Fon and Otani (1979). The notion of Pareto dominance employed says that an allocation dominates another exactly when all agents who trade contingent commodities strictly prefer their new commodity to their old one. This is a natural notion in the context of incomplete preferences where indeterminacy is interpreted in terms of sticking to the status quo: it supposes that agents who do not have strict preference for trade – either because they consider the commodity on offer not to be better than what they have, or because they do not have sufficient confidence to form determinate preferences – stick to their initial endowment.

For technical reasons, we make several assumptions on the notions of stakes and on the preferences. Since they are tangential to the general discussion in this section, they are simply noted here; statement of the relevant properties and discussion are relegated to Appendix B.

Assumption 1. For each $1 \leqslant i \leqslant n$, the stakes relation $\leqq^i$ is monotone decreasing, and the preference relation $\preceq^i$ satisfies indifference consistency and full support.

Under this assumption, we have a characterisation of Pareto optima. For a contingent commodity $x \in \Re^S_+$ and an agent i, let[16]

$$\Pi^i(x) = \left\{ \left(\frac{p(s_1)u^{i\prime}(x_{s_1})}{\sum_{t \in S} p(t)u^{i\prime}(x_t)}, \ldots, \frac{p(s_{|S|})u^{i\prime}(x_{s_{|S|}})}{\sum_{t \in S} p(t)u^{i\prime}(x_t)} \right) \mid p \in \bigcap_{z \neq x} ri(D^i(x,z)) \right\}$$

Theorem 3. *Under Assumption 1, an interior allocation $(x^1, \ldots, x^n)$ is Pareto optimal iff $\bigcap_i \Pi^i(x^i) \neq \varnothing$.*

The intuition behind this result is analogous to similar results in the literature (Rigotti and Shannon, 2005; Rigotti et al., 2008): $\Pi^i(x)$ is the set of supports of the strict upper contour set of x under $\preceq^i$, and an allocation is Pareto optimal if and only if the intersection of all such sets is non-empty. The theorem has several immediate consequences. Recall that an allocation $(x^1, \ldots, x^n)$ is a *full insurance allocation* if all the x^i are constant.

Corollary 2. *Let Assumption 1 hold, and suppose that the aggregate endowment is constant across states.*

[16] For a set A, $ri(A)$ is the relative interior of A. Note that, if A is a singleton, $ri(A) = A$.

(i) An interior full insurance allocation $(x^1, \ldots, x^n)$ *is Pareto optimal iff* $\bigcap_i \bigcap_{z \neq x^i} ri(D^i(x^i, z)) \neq \emptyset$.

(ii) If, for each i*, there exists* $\mathcal{C}^i$ *with* $\Xi^i = \{\mathcal{C}^i\}$*, then there exists an interior full insurance Pareto optimal allocation iff* $\bigcap_i ri(\mathcal{C}^i) \neq \emptyset$*. In this case, every full insurance allocation is Pareto optimal.*

The first corollary, which is a simple consequence of the fact that $\Pi^i(x^i) = \bigcap_{z \neq x^i} ri(D^i(x^i, z))$ when x^i is constant, is a general characterisation of when an interior full insurance allocation is Pareto optimal in the model proposed here. It is in the style of existing results, such as Billot et al. (2000); Rigotti and Shannon (2005); Rigotti et al. (2008). Unlike these cases, the existence of a full insurance Pareto optimal allocation does not imply that all full insurance allocations are Pareto optimal, because, in general, the relevant sets of probability measures may differ depending on the constant commodity.

The second corollary involves the special case of representation (1) where the confidence ranking is degenerate: this is essentially the unanimity model of preferences à la Bewley. The result differs slightly from that of Rigotti and Shannon (2005, Corollary 2), which also concerns the Bewley model, insofar as their result involves the intersection of the sets of probability measures of the different agents, whereas ours uses the intersections of their relative interiors. This difference is due to the fact that they take the strict preference relation as primitive, use a slightly different representation from that used here and take a stricter notion of Pareto optimality.[17]

Comparing the two corollaries brings out the relationship between the case of representation (1) and that of the unanimity model à la Bewley. Concerning the question of whether an interior full insurance allocation is Pareto optimal or not, an economy with agents represented by (1) is roughly equivalent to an economy where each agent is replaced by an agent with unanimity preferences à la Bewley, who takes as his set of probability measures the intersection of all the sets in his confidence ranking that are relevant for choices involving the commodity he is allocated.[18] *Grosso modo*, if each agent in the economy for whom more confidence is required to take decisions with higher stakes simply ignored the stakes, and always chose as if the stakes were at the lowest possible level for the commodity he

[17] They use the representation axiomatised by Bewley (2002); see Section 1.2.

[18] This is a rough statement because, for a family $\{\mathcal{C}_i | \; i \in I\}$ of closed sets, it is not necessarily the case that $\bigcap_{i \in I} ri(\mathcal{C}_i) = ri \bigcap_{i \in I} \mathcal{C}_i$.

is allocated, this would make little difference to whether the allocation is Pareto optimal or not. Although it does not follow that the set of Pareto optima would be the same as for an economy containing only agents with preferences à la Bewley, because the relevant sets of probability measures may depend on the allocation, this highlights some similarities between economies with agents à la Bewley and economies with agents who choose according to representation (1).

Things are considerably different, however, regarding the question of how fast Pareto optima can be reached. (To the extent that, as noted in the proof of Theorem 3, Pareto optima correspond to appropriately defined equilibria, this is closely related to the question of how fast the economy can arrive at equilibrium.)[19] There is often a simple, if idealised, fastest way to achieve a Pareto optimum. In particular, whenever a non-Pareto optimal allocation $(x^1, \ldots, x^n)$ is dominated by a (feasible) Pareto optimal one $(y^1, \ldots, y^n)$, then there is a "one-step" move to a Pareto optimum, which is acceptable to all agents – namely, each agent swaps x^i for y^i. (This set of "swaps" corresponds to a set of simultaneous trades between the agents.) Whenever this is the case, we say that $(y^1, \ldots, y^n)$ is *one-step accessible* from $(x^1, \ldots, x^n)$. In economies where agents have expected utility preferences, preferences à la Bewley or preferences represented by many of the standard non-expected utility theories proposed in the literature (and in particular those considered by Rigotti et al. (2008)), any Pareto dominated allocation is Pareto dominated by a Pareto optimum; in other words, any allocation has a Pareto optimum that is one-step accessible from it. This is not necessarily true for economies containing agents represented by (1), as the following example shows.

Consider a two-agent economy with two states of the world, s_1 and s_2, and suppose that each agent i has constant relative risk aversion γ^i (so the utility function is $u^i(x) = \frac{x^{1-\gamma^i}}{1-\gamma^i}$ if $\gamma^i \neq 1$ and $u^i(x) = \ln x$ if $\gamma^i = 1$). Agent 1's preferences are represented by (1), where the stakes in the choice between x and y are given by $\max_s |x(s) - y(s)|$.[20] He has

[19] Rigotti and Shannon (2005) address a related question with their notion of "equilibrium with inertia", which, approximately, is an equilibrium which Pareto dominates the initial endowment. The example below shows that, by contrast with economies whose agents have preferences à la Bewley, equilibria with inertia may not exist, even if equilibria do exist, in economies whose agents have preferences represented by (1).

[20] As noted in footnote 15, although the stakes relation is defined on contingent commodities, this yields a well-defined stakes relation on acts. It is straightforward to check that this relation satisfies the properties given in Section 2.2, and is monotone decreasing.

the following centred confidence ranking: $\{\{p \in \Delta(\Sigma)|\ 0.5 - \epsilon \leqslant p(s_1) \leqslant 0.5 + \epsilon\}|\ \epsilon \in [0, 0.45]\}$. Note that since each set in the confidence ranking is uniquely specified by an $\epsilon \in [0, 0.45]$, the cautiousness coefficient is entirely specified by a function from pairs of acts to values of ϵ. Using this formulation, the cautiousness coefficient is given by $D^1((x, y)) = \min\{\eta \max_s |x(s) - y(s)|, 0.45\}$ for $\eta > 0$, where η characterises the agent's attitude to choosing in the absence of confidence (see Section 4). Agent 2 is an expected utility decision maker with probability measure assigning 0.5 to both states.

We consider a case with no aggregate risk in the economy: the sum of allocations is w in both states. Hence allocations are of the form $((\delta_1 w, \delta_2 w), ((1 - \delta_1)w, (1 - \delta_2)w))$ for $\delta_1, \delta_2 \in [0, 1]$. It is easy to check, given Theorem 3, that the only Pareto optima are full insurance allocations. Now consider the risky endowment $(x^1, x^2) = ((\delta w, (1-\delta)w), ((1-\delta)w, \delta w))$, where $\delta \in (\frac{1}{2}, 1]$. It would seem that a natural "one-move" trade yielding a Pareto optimal allocation would be for 2 to give 1 $((\frac{1}{2} - \delta)w, (\delta - \frac{1}{2})w)$. It is easy to check that $(\frac{1}{2}w, \frac{1}{2}w) \succ^2 ((1 - \delta)w, \delta w)$. Moreover, $\sum_s p(s)u^1(\frac{1}{2}w) > \sum_s p(s)u^1(x^1_s)$ for all $p \in \bigcap_{x' \neq x^1} ri(D^1((x^1, x')))$. Were the agents to ignore the stakes and always choose as if the stakes were at their lowest level, this would be sufficient for the trade to be acceptable to both agents: that is, for the full insurance allocation $(\frac{1}{2}w, \frac{1}{2}w)$ to be one-step accessible from (x^1, x^2). However, if they take the stakes into account as specified by representation (1), there is a stronger requirement for the trade to be acceptable to agent 1, namely that $\sum_s p(s)u^1(\frac{1}{2}w) \geqslant \sum_s p(s)u^1(x^1_s)$ for all $p \in D^1((x^1, \frac{1}{2}w))$, with strict inequality for some p. By straightforward calculation, this condition holds if and only if[21]

$$(3) \qquad p \leqslant \frac{\frac{1}{2}^{1-\gamma^1} - (1 - \delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1 - \delta)^{1-\gamma^1}} \qquad \text{for all } p \in D^1((x^1, \frac{1}{2}w))$$

Hence, by the definition of D^1, $\frac{1}{2}w \not\succ^1 x^1$ whenever

$$(4) \qquad \min\{\eta w(\delta - \frac{1}{2}), 0.45\} + 0.5 > \frac{\frac{1}{2}^{1-\gamma^1} - (1 - \delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1 - \delta)^{1-\gamma^1}}$$

This inequality has solutions for various values of the parameters: it is straightforward to check, for example, that when $\gamma^1 = 2$, $\delta = \frac{3}{4}$, $w = 1500$, $\eta = 0.001$, the inequality is satisfied and so $\frac{1}{2}w \not\succ^1 x^1$. In such cases, the Pareto optimum $(\frac{1}{2}w, \frac{1}{2}w)$ is not one-step

[21] Here we consider the case where $\gamma_1 \neq 1$; the case of $\gamma_1 = 1$ can be treated similarly.

accessible from (x^1, x^2). In fact, by a similar argument, one can show that there are cases where no Pareto optimal allocation is one-step accessible from (x^1, x^2).

Proposition 3. *There may exist allocations from which no Pareto optimal allocation is one-step accessible.*

This phenomenon is basically a consequence of the dependence on stakes in representation (1), which allows agents to have determinate preferences at low stakes levels that they may withdraw at higher stakes levels. Whereas it is the former preferences – and in particular the probability measures corresponding to low levels of stakes – which determine whether an allocation is Pareto optimal or not, the latter preferences – and the associated larger sets of probability measures – determine whether an agent accepts a given trade or not. If all agents were indifferent to the stakes, and formed preferences using minimal sets of probabilities (as if the stakes were at their lowest level), then any Pareto dominated allocation would indeed be Pareto dominated by a Pareto optimal one. However, whenever there is an agent who takes account of the stakes via representation (1), he may not be confident enough in his strict preference for that Pareto optimal allocation over his initial endowment to choose the former at the appropriate level of stakes, and so sticks to the status quo. He refrains from trading, and the "one-step" move to that Pareto optimum is blocked.

We have already mentioned one interpretation of this result in terms of maximal speed of convergence. It indicates a non-trivial bound on how fast a Pareto optimal allocation can be reached: allowing any conceivable way of constructing a set of simultaneous trades (as unfeasible as it may be in practice), it may still be impossible to get to a Pareto optimum by a single set of trades if the market contains agents who incorporate confidence into their preferences, and who do not trade when they lack sufficient confidence. Another interpretation is in terms of the restrictions placed on the (theoretical) power of a social planner. In standard general equilibrium models, as well as the market under uncertainty models mentioned above, a suitably intelligent social planner who knows the agents' preferences could propose a set of simultaneous trades that would be accepted by all agents and that would bring the market to a Pareto optimum. This relies on the fact that, in these models, all allocations have a Pareto optimum that is one-step accessible. That this is not necessarily the case in the current model attests to the limited influence of such a social planner: even if he had all the information about preferences (and infinite computational power), the social planner might not be able to propose a set of simultaneous trades which leaves the economy in a Pareto optimum and is acceptable to all. The agents' tendency to demand more confidence in beliefs when the stakes are higher mean that he may not be able to persuade some of them to shift from the endowment to a Pareto optimal allocation when the stakes involved in the change are high, though they would have accepted the trade if the stakes were low. Confidence, combined with the status quo interpretation of indeterminacy of preference, can hinder (Pareto-enhancing) intervention in the market.

The natural question is, of course: how fast can a Pareto optimum be reached? Put in terms of the second interpretation offered above, this amounts to asking how many times a social planner has to intervene to bring the economy to a Pareto optimum. Let us say that a feasible allocation $(y^1, \ldots, y^n)$ is *m-step accessible* from $(x^1, \ldots, x^n)$ if there is a sequence of $m-1$ feasible allocations, the first of which Pareto dominates $(x^1, \ldots, x^n)$, the last of which is Pareto dominated by $(y^1, \ldots, y^n)$, and each of which is Pareto dominated by its successor. A Pareto optimum which is not one-step accessible may be m-step accessible: this means that, under ideal conditions, it can be reached not with a single set of trades that is acceptable to all, but rather after m consecutive sets of trades acceptable to all. If, for a given allocation, there is a Pareto optimum that is m-step accessible and none that is m'-step accessible for $m' < m$, this can be thought of as a lower bound on the how fast the economy can come to a Pareto optimum: it requires at least m sets of simultaneous trades. A social planner has to intervene at least m times. There is, however, no general lower bound that applies to all allocations in all economies.

Proposition 4. *There may exist allocations from which no Pareto optimal allocation is m-step accessible, for any finite m.*

When there are agents whose preferences incorporate their confidence in beliefs, and who stick to the status quo when they do not have enough confidence to take a choice, it may thus be theoretically impossible for the market to arrive at a Pareto optimal allocation in finite time. Because, quite simply, there may not exist a finite sequence of sets of trades, where all agents have sufficient confidence to accept the trades and where the sequence reaches a Pareto optimum. Confidence, combined with taking the status quo option – and not trading – when one is not sufficiently confident in any option, adds considerable friction into the economy.

7 Conclusion

People may be more or less confident in their beliefs; moreover, they may have incomplete preferences. In this paper, a theory which relates incompleteness of preferences to confidence in beliefs was proposed. It is based on the following maxim: one has a determinate preference over a pair of acts if and only if one's confidence in the beliefs needed to form the preference matches up to the stakes involved in the choice between the acts. In the absence of sufficient confidence, preferences are indeterminate.

A formal decision rule conforming to this maxim was proposed. The decision maker's confidence in his beliefs is modelled by a confidence ranking – a nested family of sets of probability measures. A cautiousness coefficient assigns to any decision a level of confidence relevant for that decision (represented formally by a set in the confidence ranking), which is determined by the stakes involved. The decision rule according to which one act is preferred to another if it has higher expected utility according to all the probability measures in the appropriate set was axiomatised. Moreover, comparative statics analysis of the relative decisiveness of decision makers, as well as of their confidence in preferences, was undertaken, and the cautiousness coefficient was seen to correspond to the decision maker's attitude to choosing in the absence of confidence.

Under one possible interpretation, indeterminate preference can be understood in terms of deferral, be it to someone else, to one's later self, or to a status quo. In order to consider the consequences of the theory in situations where deferral is not an option, a preference relation was introduced representing preferences in the absence of a deferral option. An axiomatic account of the relationship between the decision maker's preferences in the presence and absence of a deferral option was proposed.

Finally, possible consequences of the model in a market setting were considered, where indeterminacy of preferences was interpreted in terms of refusal to trade. On the one hand, a characterisation of Pareto optima was provided, and the relationship to analogous results for other decision models was discussed. On the other hand, it was shown that there may exist Pareto dominated allocations that are not dominated by any Pareto optimum. It follows that, in markets with agents incorporating confidence into their preferences and refusing to trade when they are not confident enough to form a determinate preference, it may be theoretically impossible for the market to come to a Pareto optimum by a single set of trades accepted by all. Moreover, there are cases where no finite sequence of sets of trades, accepted by all, can bring the market to a Pareto optimum.

A Proofs

Throughout the Appendix, B will denote the space of all real-valued functions on S, and $ba(S)$ will denote the set of finitely additive real-valued set functions on S, both under the Euclidean topology. Recall that, under this topology, $ba(S)$ is locally convex (Aliprantis and Border, 2007, §5.12). B is equipped with the standard order: $a \leqslant b$ iff $a(s) \leqslant b(s)$ for all $s \in S$. For $x \in \Re$, we define x^* to be the constant function taking value x.

A.1 Proof of Theorem 1

The main part of the result is to show the sufficiency of the axioms for the representation (direction (i) to (ii)), the proof of which proceeds as follows. By standard arguments, we obtain a von Neumann-Morgernstern utility function on the consequences, which allows us to work with real-valued functions on S instead of acts. For each non-minimal non-trivial stakes level r, we define a preference relation $\preceq_r$ on these functions, which can be thought of as representing the preferences between corresponding acts considered "as if" the choices had stakes r. We show (Lemma 5) that each $\preceq_r$ is a non-trivial, monotonic, affine, Archimedean pre-order, whence, by Gilboa et al. (2010, Corollary 1) (which is a version of a Ghirardato et al. (2004, Proposition A.2)), there is a closed convex set of probability measures $\mathcal{C}_r$ representing $\preceq_r$ according to the unanimity rule. Moreover, Lemma 10 shows that, whenever there is a minimal stakes level and non-trivial choices (ie. choices other than between an act and itself) that have this stakes level, the preference relation for choices with these stakes can be represented according to the unanimity rule with the intersection of the $\mathcal{C}_r$ for the other stakes levels. By Lemma 7, the $\mathcal{C}_r$ form a nested family of sets, and we thus have a confidence ranking. By Lemmas 8 and 9, this confidence ranking is continuous. By construction, the function that assigns to any stakes level r the set $\mathcal{C}_r$ is a well-defined cautiousness coefficient.

Now we proceed with the proof. First we assume (i); we will show (ii). We require the following lemmas.

Lemma 1. *There exists a non-constant utility function u representing the restriction of $\preceq$ to the constant acts.*

Proof. By A1, A2, A5, A4 and A7, the restriction of $\preceq$ to constant acts is non-trivial complete, reflexive and satisfies independence and continuity. Moreover, it can be shown to be transitive. For any $c, c', c'' \in \Delta(X)$, suppose that $c \preceq c'$ and $c' \preceq c''$. First consider the case where $c \nsim c'$ and $c' \nsim c''$. By richness of $\leqq$, there exist $\alpha \in (0,1]$ and $d \in \Delta(X)$ such that $(c_\alpha d, c''_\alpha d) \leqq (c, c')$ and $(c_\alpha d, c''_\alpha d) \leqq (c', c'')$. By A3, $c_\alpha d \preceq c''_\alpha d$, from which it follows by A1 and A4, that $c \preceq c''$, as required. If $c \sim c'$ and $c' \nsim c''$, then by richness of $\leqq$, there exists $\alpha \in (0,1]$ and $d \in \Delta(X)$ such that $(c_\alpha d, c''_\alpha d) \leqq (c', c'')$; by A3, $c_\alpha d \preceq c''_\alpha d$, and so by A1 and A4, $c \preceq c''$. The case where $c' \sim c''$ and $c \nsim c'$ is treated similarly. The case where $c \sim c' \sim c''$ follows directly from A3. The existence of u follows from the standard von Neumann-Morgenstern theorem. □

Let $K = u(\Delta(X))$ and $B(K)$ be the set of functions in B taking values in K. There is thus a many-to-one mapping between acts in $\mathcal{A}$ and elements of $B(K)$, given by $a = u \circ f$, $f \in \mathcal{A}$. With slight abuse of notation, we use $\preceq$ to denote the order generated on $B(K)$ by $\preceq$ under this mapping, and $\leqq$ to denote the order generated on $B(K) \times B(K)$ by $\leqq$. (That $\preceq$ and $\leqq$ are well-defined on $B(K)$ follows from A5 and the extensionality of $\leqq$ respectively.)

Lemma 2. *For each $a, a', b, b' \in B(K)$ with $a \neq b$ and $a' \neq b'$, there exists $\alpha \in (0,1]$ and $l \in B(K)$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \equiv (a', b')$.*

Proof. If $(a, b) \equiv (a', b')$, then there is nothing to show. Suppose without loss of generality that $(a, b) < (a', b')$; the other case is treated similarly. By richness of $\leqq$, there exist $\beta \in (0,1]$ and $l \in \mathcal{A}$ such that $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \geqq (a', b')$. If $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \equiv (a', b')$, then the result has been established; if not, then by continuity of $\leqq$, there exists $\alpha \in (\beta, 1)$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \equiv (a', b')$, as required. □

Let $\mathcal{S}$ be the set of equivalence classes of $\leqq$. As standard, $\leqq$ on $B(K) \times B(K)$ generates a relation on $\mathcal{S}$, which will be denoted $\leqslant$ (with symmetric and asymmetric components $=$ and $<$ respectively): for $r, s \in \mathcal{S}$, $r \leqslant s$ iff, for any $(f, g) \in r$ and $(f', g') \in s$, $(f, g) \leqq (f', g')$. $r \in \mathcal{S}$ is a minimal (respectively maximal) element if $r \leqslant s$ (resp. $r \geqslant s$) for all $s \in \mathcal{S}$. Note that, since $\leqslant$ is a linear ordering, there is at most one minimal (resp. maximal) element; if it exists, we denote the minimal (resp. maximal) element by $\underline{\mathcal{S}}$ (resp. $\overline{\mathcal{S}}$). $r \in \mathcal{S}$ is non-trivial if there exists $(a, b) \in r$ with $a \neq b$. It follows from the continuity of $\leqq$

that every non-minimal and non-maximal element in $\mathcal{S}$ is non-trivial. Let $\mathcal{S}^{nt}$ be the set of non-trivial elements in $\mathcal{S}$, and let $\mathcal{S}^{+}$ be the set of non-minimal non-trivial elements. For each $r \in \mathcal{S}^{nt}$, let $\preceq_r$ be the reflexive binary relation on $B(K)$ such that, for all $a, b \in B(K)$ with $a \neq b$, $a \preceq_r b$ iff there exists $l \in B(K)$ and $\alpha \in (0,1]$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$ and $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$.

Lemma 3. *For each $r \in \mathcal{S}^{nt}$, $a, b, l \in B(K)$, and $\alpha \in (0,1]$, if $(\alpha a+(1-\alpha)l, \alpha b+(1-\alpha)l) \in r$, then $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ iff $a \preceq_r b$.*

Proof. The result is immediate if $a = b$, so consider $a, b \in B(K)$ with $a \neq b$. It needs to be shown that, whenever there exists $l, m \in B(K)$ and $\alpha, \beta \in (0,1]$ such that $(\alpha a + (1-\alpha)l, \alpha b+(1-\alpha)l) \equiv (\beta a+(1-\beta)m, \beta b+(1-\beta)m)$, then $\alpha a+(1-\alpha)l \preceq \alpha b+(1-\alpha)l$ iff $\beta a+(1-\beta)m \preceq \beta b+(1-\beta)m$. Without loss of generality, suppose that $\beta \leqslant \alpha$. Consider first the case where $\beta < \alpha$. Note that $\beta a+(1-\beta)m = \frac{\beta}{\alpha}(\alpha a+(1-\alpha)l)+(1-\frac{\beta}{\alpha})(\frac{\alpha\beta-\beta}{\alpha-\beta}l+\frac{\alpha-\alpha\beta}{\alpha-\beta}m)$, where $\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m \in B(K)$, since it is a $\frac{\alpha\beta-\beta}{\alpha-\beta}$-mix of l and m; similarly for $\beta b+(1-\beta)m$. Let $f, g, h \in \mathcal{A}$ be such that $\alpha a+(1-\alpha)l = u\circ f$, $\alpha b+(1-\alpha)l = u \circ g$ and $\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m = u \circ h$; so $\beta a+(1-\beta)m = u\circ f_{\frac{\beta}{\alpha}}h$ and $\beta b+(1-\beta)m = u \circ g_{\frac{\beta}{\alpha}}c$. Since $(f, g) \equiv (f_{\frac{\beta}{\alpha}}h, g_{\frac{\beta}{\alpha}}h)$, by A6, $f' \asymp g'$ iff $f_{\frac{\beta}{\alpha}}h \asymp g_{\frac{\beta}{\alpha}}h$. Hence, by A4, $f \preceq g$ iff $f_{\frac{\beta}{\alpha}}h \preceq g_{\frac{\beta}{\alpha}}h$. So $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ iff $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$, as required.

Now consider the case where $\beta = \alpha$. If $l = m$, the result is immediate, so suppose that $l \neq m$. Note that if there exists $\epsilon \in (0,1]$ and $k \in B(K)$ with $\epsilon \neq \alpha$ and $(\epsilon a+(1-\epsilon)k, \epsilon b+(1-\epsilon)k) \equiv (\alpha a+(1-\alpha)l, \alpha b+(1-\alpha)l)$, then, by applying the reasoning in the case above, we get $\alpha a+(1-\alpha)l \preceq \alpha b+(1-\alpha)l$ iff $\epsilon a+(1-\epsilon)k \preceq \epsilon b+(1-\epsilon)k$ iff $\beta a+(1-\beta)m \preceq \beta b + (1-\beta)m$, as required. This obviously holds if all pairs of non-identical acts have the same stakes, so suppose that this is not the case. By Lemma 2, there exists $n \in B(K)$ and $\gamma \in (0,1)$ such that $(\gamma(\alpha a + (1-\alpha)l) + (1-\gamma)n, \gamma(\alpha b + (1-\alpha)l) + (1-\gamma)n) \not\equiv (\alpha a+(1-\alpha)l, \alpha b+(1-\alpha)l)$. By Lemma 2 again, there exists $\delta \in (0,1)$ and $n' \in B(K)$ such that $(\delta(\gamma(\alpha a+(1-\alpha)l)+(1-\gamma)n)+(1-\delta)n', \delta(\gamma(\alpha b+(1-\alpha)l)+(1-\gamma)n)+(1-\delta)n') \equiv (\alpha a+(1-\alpha)l, \alpha b+(1-\alpha)l) \in r$. However, $\delta(\gamma(\alpha a+(1-\alpha)l)+(1-\gamma)n)+(1-\delta)n' = \alpha\gamma\delta a+(1-\alpha\gamma\delta)(\frac{\delta-\alpha\gamma\delta}{1-\alpha\gamma\delta}(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}l+\frac{1-\gamma}{1-\alpha\gamma}n)+\frac{1-\delta}{1-\alpha\gamma\delta}n')$, with $\frac{\delta-\alpha\gamma\delta}{1-\alpha\gamma\delta}(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}l+\frac{1-\gamma}{1-\alpha\gamma}n)+\frac{1-\delta}{1-\alpha\gamma\delta}n' \in B(K)$ since it is a mix of elements of $B(K)$. So $\alpha\gamma\delta$ and $\frac{\delta-\alpha\gamma\delta}{1-\alpha\gamma\delta}(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}l + \frac{1-\gamma}{1-\alpha\gamma}n) + \frac{1-\delta}{1-\alpha\gamma\delta}n'$ have the properties required above, and the result is established.

□

Lemma 4. *For all $r, s \in \mathcal{S}^{nt}$ with $r \geqslant s$, $\preceq_r \subseteq \preceq_s$.*

Proof. If $s = r$, there is nothing to show, so suppose not and consider $a, b \in B$ such that $a \preceq_r b$. If $a = b$, the result follows from the reflexivity of $\preceq_s$ and $\preceq_r$; henceforth suppose that this is not the case. Without loss of generality, it can be assumed that $(a, b) \in r$. (If not, replace a, b with $\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l$ where $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$ and continue as below.) It follows from Lemma 3 that $a \preceq b$. Let $\beta \in (0, 1)$ and $m \in B(K)$, be such that $(\beta a + (1-\beta)m, \beta b + (1-\beta)m) \in s$ (such β and m exist by Lemma 2). A6 and A4 imply that $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$, and hence $a \preceq_s b$, as required. □

Recall that a binary relation $\preceq$ on $B(K)$ is

- *non-trivial* if there exists $a, b \in B(K)$ such that $a \preceq b$ but not $a \succeq b$.
- *monotonic* if, for all $a, b, \in B(K)$, if $a \leqslant b$ then $a \preceq b$.
- *affine* if, for all $a, b, c \in B(K)$ and $\alpha \in (0, 1)$, $a \preceq b$ iff $\alpha a + (1-\alpha)c \preceq \alpha b + (1-\alpha)c$.
- *Archimedean* if, for all $a, b, c \in B(K)$, the sets $\{\alpha \in [0, 1] \mid \alpha a + (1-\alpha)b \succeq c\}$ and $\{\alpha \in [0, 1] \mid \alpha a + (1-\alpha)b \preceq c\}$ are closed in $[0, 1]$.
- a *pre-order* if $\preceq$ is reflexive and transitive.

Lemma 5. *For every $r \in \mathcal{S}^{nt}$, $\preceq_r$ is a non-trivial, monotonic, affine pre-order. Moreover, if $r \in \mathcal{S}^+$, $\preceq_r$ is Archimedean.*

Proof. Non-triviality. By A2, $\preceq$ is non-trivial; by A5 and A1, it follows that the restriction of $\preceq$ to $\Delta(X)$ is non-trivial. But $\preceq_r$ coincides with $\preceq$ on $\Delta(X)$, so it is non-trivial.

Monotonicity. Suppose that $a \leqslant b$ and $a \neq b$ (the result is immediate for $a = b$). Then, $\alpha a + (1-\alpha)l \leqslant \alpha b + (1-\alpha)l$ for $l \in B(K)$ and $\alpha \in (0, 1]$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$. By monotonicity (A5), $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$, and so $a \preceq_r b$.

Affineness. The result is immediate if $a = b$; henceforth suppose not. By Lemma 2, there exists $\beta \in (0, 1]$ and $l \in B(K)$ such that $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \in r$. Consider $\beta(\alpha a + (1-\alpha)c) + (1-\beta)l$ and $\beta(\alpha b + (1-\alpha)c) + (1-\beta)l$: by Lemma 2, there exists $\gamma \in (0, 1]$ and $m \in B(K)$ such that $(\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m, \gamma(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m) \in r$. Note that $\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m = \alpha\gamma(\beta a + (1-\beta)l) + (1-\alpha\gamma)(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m)$, where $\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m \in B(K)$ since it is a mix of elements of $B(K)$.

Similarly for b. Let $f, g, h \in \mathcal{A}$ be such that $\beta a + (1-\beta)l = u \circ f$, $\beta b + (1-\beta)l = u \circ g$ and $\frac{\gamma - \alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m = u \circ h$. Since $(f, g) \equiv (f_{\alpha\gamma}h, g_{\alpha\gamma}h)$, by A6 and A4, $\beta a + (1-\beta)l \leq \beta b + (1-\beta)l$ iff $\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m \leq \gamma(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m$. But since $\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m = \beta\gamma(\alpha a + (1-\alpha)c) + (1-\beta\gamma)(\frac{\gamma-\beta\gamma}{1-\beta\gamma}l + \frac{1-\gamma}{1-\beta\gamma}m)$, and similarly for b, it follows that $a \preceq_r b$ iff $\alpha a + (1-\alpha)c \preceq_r \alpha b + (1-\alpha)c$, as required.

Pre-order. Reflexivity follows from the definition of $\preceq_r$. As for transitivity, suppose that $a \preceq_r b$ and $b \preceq_r c$ and that $a \neq b \neq c$ (if $a = b$, $b = c$ or $a = c$, the result is immediate). By Lemma 2, there exists $l \in B(K)$ and $\alpha \in (0, 1]$ such that $(\alpha a + (1-\alpha)l, \alpha c + (1-\alpha)l) \in r$. Moreover, there exists $m, n \in B(K)$ and $\beta, \gamma \in (0, 1]$ such that $(\beta(\alpha a + (1-\alpha)l) + (1-\beta)m), \beta(\alpha b + (1-\alpha)l) + (1-\beta)m)) \in r$ and $(\gamma(\alpha b + (1-\alpha)l) + (1-\gamma)m), \gamma(\alpha c + (1-\alpha)l) + (1-\gamma)m) \in r$. Since $a \preceq_r b$ and $b \preceq_r c$, $\beta(\alpha a + (1-\alpha)l) + (1-\beta)m \preceq \beta(\alpha b + (1-\alpha)l) + (1-\beta)m)$ and $\gamma(\alpha b + (1-\alpha)l) + (1-\gamma)m \preceq \gamma(\alpha c + (1-\alpha)l) + (1-\gamma)m$. Hence, by A3, $\alpha a + (1-\alpha)l \preceq \alpha c + (1-\alpha)l$, and so $a \preceq_r c$, as required.

Archimedean. Let r be a non-minimal element of $\mathcal{S}$. Consider $\{\alpha \in [0, 1] \mid \alpha a + (1-\alpha)b \succeq_r c\}$; the other case is dealt with similarly. Let $\bar{\alpha}$ be a limit point of this set, and without loss of generality, assume that $(\bar{\alpha} a + (1-\bar{\alpha})b, c) \in r$ (if not, replace a, b, c with appropriate mixtures for which this is the case). It needs to be shown that $\bar{\alpha} a + (1-\bar{\alpha})b \succeq_r c$. If $\bar{\alpha} a + (1-\bar{\alpha})b = c$ the result is immediate; suppose henceforth that this is not the case. If there is a open interval I in $\{\alpha \in [0, 1] \mid \alpha a + (1-\alpha)b \succeq_r c\}$ such that $\bar{\alpha}$ is a limit point of I and such that $(\beta a + (1-\beta)b, c) \leqq (\bar{\alpha} a + (1-\bar{\alpha})b, c)$ for all $\beta \in I$, then, by Lemma 4, $\beta a + (1-\beta)b \succeq c$ for all $\beta \in I$, whence $\bar{\alpha} a + (1-\bar{\alpha})b \succeq c$ by A7, and so $\bar{\alpha} a + (1-\bar{\alpha})b \succeq_r c$ as required.

Now suppose that there is no such interval. Since r is a non-minimal element of $\mathcal{S}$, by the continuity of $\leqq$ and Lemma 2, there exists $l \in B(K)$ and $\bar{\delta} \in (0, 1)$ such that $(\bar{\delta}(\bar{\alpha} a + (1-\bar{\alpha})b) + (1-\bar{\delta})l, \bar{\delta} c + (1-\bar{\delta})l) < (\bar{\alpha} a + (1-\bar{\alpha})b, c)$. Let $\gamma = \min\{\delta \in (\bar{\delta}, 1] | \, (\delta(\bar{\alpha} a + (1-\bar{\alpha})b) + (1-\delta)l, \delta c + (1-\delta)l) \geqq (\bar{\alpha} a + (1-\bar{\alpha})b, c)\}$ (by continuity of $\leqq$ this is a minimum). Consider any $\delta \in (\bar{\delta}, \gamma)$; by the definition of γ, $(\delta(\bar{\alpha} a + (1-\bar{\alpha})b) + (1-\delta)l, \delta c + (1-\delta)l) < (\bar{\alpha} a + (1-\bar{\alpha})b, c)$. Note moreover that $\delta(\bar{\alpha} a + (1-\bar{\alpha})b) + (1-\delta)l = \bar{\alpha}(\delta a + (1-\delta)l) + (1-\bar{\alpha})(\delta b + (1-\delta)l)$. So, by the continuity of $\leqq$, there is an open interval $I_\delta \subseteq (0, 1)$ containing $\bar{\alpha}$ such that, for all $\beta \in I_\delta$, $(\beta(\delta a + (1-\delta)l) + (1-\beta)(\delta b + (1-$

$\delta)l), \delta c + (1-\delta)l) \leqq (\bar{\alpha}a + (1-\bar{\alpha})b, c)$. Note that $I_\delta \cap \{\alpha \in [0,1] \mid \alpha a + (1-\alpha)b \succeq_r c\}$ is non-empty, since $\bar{\alpha}$ is a limit point of $\{\alpha \in [0,1] \mid \alpha a + (1-\alpha)b \succeq_r c\}$. Furthermore, since $\beta(\delta a + (1-\delta)l) + (1-\beta)(\delta b + (1-\delta)l) = \delta(\beta a + (1-\beta)b) + (1-\delta)l$, Lemma 4 implies that $\beta(\delta a + (1-\delta)l) + (1-\beta)(\delta b + (1-\delta)l) \succeq \delta c + (1-\delta)l$ for all $\beta \in I_\delta \cap \{\alpha \in [0,1] \mid \alpha a + (1-\alpha)b \succeq_r c\}$. It follows by A7 that $\bar{\alpha}(\delta a + (1-\delta)l) + (1-\bar{\alpha})(\delta b + (1-\delta)l) \succeq \delta c + (1-\delta)l$. Since this holds for all $\delta \in (\bar{\delta}, \gamma)$, it follows by A7 that $\gamma(\bar{\alpha}a + (1-\bar{\alpha})b) + (1-\gamma)l \succeq \gamma c + (1-\gamma)l$; whence, since $(\gamma(\bar{\alpha}a + (1-\bar{\alpha})b) + (1-\gamma)l, \gamma c + (1-\gamma)l) \in r$, $\bar{\alpha}a + (1-\bar{\alpha})b \succeq_r c$, as required.

□

Lemma 6. *For each $r \in \mathcal{S}^+$, there exists a unique closed convex set of probabilities $\mathcal{C}_r$ representing $\preceq_r$ according to the following equation: for all $a, b \in B$, $a \preceq_r b$ iff*

$$\sum_{s \in S} a(s)p(s) \leqslant \sum_{s \in S} b(s)p(s) \qquad \textit{for all } p \in \mathcal{C}_r \tag{5}$$

Proof. This follows from Lemma 5, by Gilboa et al. (2010, Corollary 1),[22] which establishes such a representation for non-trivial, monotonic, affine, Archimedean pre-orders. □

Lemma 7. *For all $r, s \in \mathcal{S}^+$ with $r \geqslant s$, $\mathcal{C}_s \subseteq \mathcal{C}_r$.*

Proof. This follows directly from Lemma 4 and Ghirardato et al. (2004, Proposition A.1).

□

Lemma 8. *For all $r \in \mathcal{S}^+$, $\mathcal{C}_r = \overline{\bigcup_{r' < r} \mathcal{C}_{r'}}$.*

Proof. By Lemma 7, $\mathcal{C}_r \supseteq \mathcal{C}_{r'}$ for all $r' < r$. Suppose, for reductio, that $\mathcal{C}_r \supsetneq \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$, so that there exists a point (probability measure) $p \in \mathcal{C}_r \backslash \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. By a separating hyperplane theorem (Aliprantis and Border (2007, 5.80)), there is a continuous linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. Since S is finite (so B is finite-dimensional), B is reflexive, and, by the standard isomorphism between $ba(S)$ and B^*, it follows that $ba(S)^*$ is isometrically isomorphic to B (Dunford and Schwartz, 1958, IV.3); hence there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality, α, ϕ and a can be chosen so that $\alpha \in K$, $a \in B(K)$, $(a, \alpha^*) \in r$. By Lemma 2, there exists $l \in B(K)$ and $\beta \in (0,1)$ such that

[22] See Ghirardato et al. (2004, Proposition A.2) for a related result.

$(\beta a + (1-\beta)l, \beta\alpha^* + (1-\beta)l) < (a, \alpha^*)$. Let $\beta' = \min\{\gamma \in [\beta, 1] | (\gamma a + (1-\gamma)l, \gamma\alpha^* + (1-\gamma)l) \geqq (a, \alpha^*)\}$ (this is a minimum by the continuity of $\leqq$). Taking $f, h \in \mathcal{A}$ and $c \in \Delta(X)$ such that $u \circ f = a$, $u \circ h = l$ and $u(c) = \alpha$, it follows, by the construction, that for any $\gamma \in (\beta, \beta')$, $c_\gamma h \preceq f_\gamma h$. However, by construction, $c_{\beta'} h \npreceq f_{\beta'} h$, contradicting A7. Hence $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$.

□

Lemma 9. *For all non-maximal $r \in \mathcal{S}^+$, $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.*

Proof. By Lemma 7, $\mathcal{C}_r \subseteq \mathcal{C}_{r'}$ for all $r' > r$. Suppose, for reductio, that $\mathcal{C}_r \subsetneq \bigcap_{r'>r} \mathcal{C}_{r'}$, so that there exists a point (probability measure) $p \in \bigcap_{r'>r} \mathcal{C}_{r'} \backslash \mathcal{C}_r$. By a separating hyperplane theorem (Aliprantis and Border (2007, 5.80)), there is a continuous linear functional ϕ on $ba(S)$, an $\alpha \in \Re$ and an $\epsilon > 0$ such that $\phi(p) \leqslant \alpha - \epsilon$ and $\alpha \leqslant \phi(q)$ for all $q \in \mathcal{C}_r$. As in the proof of Lemma 8, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s\in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality, α, ϕ and a can be chosen so that $\alpha \in K$, $a \in B(K)$, $(a, \alpha^*) \in r$. By Lemma 2, there exists $l \in B(K)$ and $\beta \in (0, 1)$ such that $(\beta a + (1-\beta)l, \beta\alpha^* + (1-\beta)l) > (a, \alpha^*)$. Let $\beta' = \min\{\gamma \in [\beta, 1] | (\gamma a + (1-\gamma)l, \gamma\alpha^* + (1-\gamma)l) \leqq (a, \alpha^*)\}$. Taking $f, h \in \mathcal{A}$ and $c, d \in \Delta(X)$ such that $u \circ f = a$, $u \circ h = l$, $u \circ c = \alpha^*$ and $u \circ d = (\alpha - \frac{\epsilon}{2})^*$, it follows, by the construction, that $d_{\beta'} h \prec c_{\beta'} h \preceq f_{\beta'} h$ but that $d_\gamma h \npreceq f_\gamma h$ for all $\gamma \in (\beta, \beta')$. By continuity of $\leqq$ and Lemmas 3 and 4, it follows that for any $h' \in \mathcal{A}$ and $\gamma' \in (0, 1)$ such that $(d_{\gamma'} h', f_{\gamma'} h') > (d, f)$, $d_{\gamma'} h' \npreceq f_{\gamma'} h'$, contradicting A8; hence $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.

□

Lemma 10. *Let $\preceq_{\bigcap \mathcal{S}}$ be the relation on $B(K)$ generated by (5) with the set of probability measures $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$. If $\underline{S}$ is a minimal non-trivial element of $\mathcal{S}$, then $\preceq_{\underline{S}} = \preceq_{\bigcap \mathcal{S}}$.*

Proof. By Lemma 4, $\preceq_{\underline{S}} \supseteq \bigcup_{r\in\mathcal{S}^+} \preceq_r$, and so $\preceq_{\underline{S}} \supseteq \preceq_{\bigcap \mathcal{S}}$. For the inverse containment, suppose that $a \preceq_{\underline{S}} b$. Since $\preceq_{\underline{S}}$ is an affine pre-order (by Lemma 5), it follows that $\alpha(b - a) + x^* \succeq_{\underline{S}} x^*$, for any $\alpha > 0$, $x \in \Re$ such that $\alpha(b - a) + x^*, x^* \in B(K)$. Choose α and x such that this is the case. If x is minimal in K then $\alpha(b - a) + x^* \succeq_{\bigcap \mathcal{S}} x^*$ since $\preceq_{\bigcap \mathcal{S}}$ is monotonic, so suppose not. By A8, for any $\epsilon > 0$, there exists $s > \underline{S}$ such that $\alpha(b - a) + x^* \succeq_s (x - \epsilon)^*$; therefore $\alpha(b - a) + x^* \succeq_{\bigcap \mathcal{S}} (x - \epsilon)^*$ for all $\epsilon > 0$. Since $\preceq_{\bigcap \mathcal{S}}$ is Archimedean, it follows that $\alpha(b - a) + x^* \succeq_{\bigcap \mathcal{S}} x^*$, and, by the affineness of $\preceq_{\bigcap \mathcal{S}}$, $a \preceq_{\bigcap \mathcal{S}} b$, as required.

□

Conclusion of the proof of Theorem 1. Define

$$\Xi = \begin{cases} \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} & \text{if } \mathcal{S} = \mathcal{S}^+ \\ \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \{\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r\} & \text{if } \mathcal{S} = \mathcal{S}^+ \cup \{\underline{\mathcal{S}}\} \\ \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \{\overline{\bigcup_{r \in \mathcal{S}^+} \mathcal{C}_r}\} & \text{if } \mathcal{S} = \mathcal{S}^+ \cup \{\overline{\mathcal{S}}\} \\ \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \{\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r, \overline{\bigcup_{r \in \mathcal{S}^+} \mathcal{C}_r}\} & \text{if } \mathcal{S} = \mathcal{S}^+ \cup \{\underline{\mathcal{S}}, \overline{\mathcal{S}}\} \end{cases}$$

where the $\mathcal{C}_r$ are as specified in Lemma 6 and where, in the last two cases, $\overline{\mathcal{S}}$ is understood to be a maximal element of $\mathcal{S}$ not belonging to $\mathcal{S}^+$. It follows from Lemma 7 that Ξ is a nested family of sets. Since the $\mathcal{C}_r$ are closed and convex for all $r \in \mathcal{S}^+$ (Lemma 6), Ξ is a confidence ranking. By Lemmas 8 and 9, Ξ is continuous. D is defined as follows: for all $(f,g) \in \mathcal{A} \times \mathcal{A}$, if $[(f,g)] \in \mathcal{S}^+$, then $D((f,g)) = \mathcal{C}_{[(f,g)]}$, if $(f,g) \in \underline{\mathcal{S}}$, then $D((f,g)) = \bigcap_{s \in \mathcal{S}^+} \mathcal{C}_s$, and if $(f,g) \in \overline{\mathcal{S}}$, then $D((f,g)) = \overline{\bigcup_{s \in \mathcal{S}^+} \mathcal{C}_s}$. Order preservation and surjectivity of D are immediate from the definition and Lemma 7. By construction and Lemma 10, u, Ξ, D represent $\preceq$ according to (1).

Now consider the clause concerning the centering axiom A9. Let $\preceq_{\bigcap \mathcal{S}}$ be the relation on $B(K)$ generated by (5) with the set of probability measures $\bigcap_{r \in \mathcal{S}^{nt}} \mathcal{C}_r$. Obviously $\preceq_{\bigcap \mathcal{S}} \supseteq \bigcup_{r \in \mathcal{S}^{nt}} \preceq_r$; by A9, the latter relation is complete, and hence so is $\preceq_{\bigcap \mathcal{S}}$. It follows from the form of (5) that $\bigcap_{r \in \mathcal{S}^{nt}} \mathcal{C}_r$ is a singleton, as required.

The direction from (ii) to (i) is generally straightforward. The only interesting case is continuity (A7). Consider any $f, g, h \in \mathcal{A}$, and the set $\{(\alpha, \beta) \in [0,1]^2 \mid f_\alpha h \preceq g_\beta h\}$. Suppose that (α^*, β^*) is a limit point of this set, and consider a sequence $((\alpha_i, \beta_i))$ of members of the set with $(\alpha_i, \beta_i) \to (\alpha^*, \beta^*)$. If there exists a subsequence of $((\alpha_i, \beta_i))$, tending to (α^*, β^*), such that $(f_{\alpha_{i_n}} h, g_{\beta_{i_n}} h) \geqq (f_{\alpha^*} h, g_{\beta^*} h)$ for all $(\alpha_{i_n}, \beta_{i_n})$ in this sequence, then the case can be treated in a similar way to that used below, relying on the continuity of the confidence ranking. Henceforth, we consider the case where there is no such sequence. In this case, there exists $f', g' \in \mathcal{A}$ with $(f', g') < (f_{\alpha^*} h, g_{\beta^*} h)$. Moreover, by the continuity of $\leqq$, for each such (f', g'), there is an open interval around (α^*, β^*) such that $(f_\gamma h, g_\delta h) \geqq (f', g')$ for any (γ, δ) in this interval. Hence, for each such (f', g'), there is a subsequence $((\alpha_{j_n^{[(f',g')]}}, \beta_{j_n^{[(f',g')]}}))$ of $((\alpha_i, \beta_i))$, tending to (α^*, β^*), with $(f_{\alpha_{j_n^{[(f',g')]}}} h, g_{\beta_{j_n^{[(f',g')]}}} h) \geqq (f', g')$ for all $n \in \mathbb{N}$. It follows, since D is order-preserving, that for all $n \in \mathbb{N}$, $\sum_{s \in S} u(f_{\alpha_{j_n^{[(f',g')]}}} h(s)).p(s) \leqslant \sum_{s \in S} u(g_{\beta_{j_n^{[(f',g')]}}} h(s)).p(s)$, for all $p \in D((f',g'))$. Hence, by the continuity of the linear representation, it follows that $\sum_{s\in S} u(f_{\alpha^*}h(s)).p(s) \leqslant \sum_{s\in S} u(g_{\beta^*}h(s)).p(s))$, for all $p \in D((f',g'))$. Since this holds for every $(f',g') < (f,g)$, and since, by the continuity of the confidence ranking and the surjectivity of D, $D((f,g)) = \overline{\bigcup_{(f',g')<(f,g)} D((f',g'))}$, there cannot be a $q \in D((f,g))$ such that $\sum_{s\in S} u(f_{\alpha^*}h(s)).q(s) > \sum_{s\in S} u(g_{\beta^*}h(s)).q(s)$. So $f_{\alpha^*}h \preceq g_{\beta^*}h$, and hence $\{(\alpha,\beta) \in [0,1]^2 | \ f_\alpha h \preceq g_\beta h\}$ is closed, as required.

Finally, consider the uniqueness clause. Uniqueness of u follows from the von Neumann-Morgenstern theorem. As regards uniqueness of Ξ, proceed by reductio; suppose that u, Ξ_1, D_1 and u, Ξ_2, D_2 both represent $\preceq$ according to (1), with $\Xi_1 \neq \Xi_2$. So, for some $f,g \in \mathcal{A}$, $D_1((f,g)) \neq D_2((f,g))$ (by the surjectivity of the D_i). Suppose, without loss of generality, that $p \in D_1((f,g)) \backslash D_2((f,g))$. By a separating hyperplane theorem (Aliprantis and Border, 2007, 5.80), there is a continuous linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in D_2((f,g))$. Since S is finite (so B is finite-dimensional), B is reflexive, and, by the standard isomorphism between $ba(S)$ and B^*, it follows that $ba(S)^*$ is isometrically isomorphic to B (Dunford and Schwartz, 1958, IV.3); hence there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s\in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality ϕ, a and α can be chosen so that $\alpha \in K$, $a \in B(K)$ and $(a, \alpha^*) \equiv (u \circ f, u \circ g)$. Taking $h \in \mathcal{A}$ such that $u \circ h = a$ and $c \in \Delta(X)$ such that $u(c) = \alpha$, we have that $\sum_{s\in S} u(h(s))p(s) \geqslant \sum_{s\in S} u(c)p(s)$ for all p s.t. $p \in D_2((h,c))$, whereas this is not the case for all p s.t. $p \in D_1((h,c))$, contradicting the assumption that both u, Ξ_1, D_1 and u, Ξ_2, D_2 represent $\preceq$. A similar argument establishes the uniqueness of D.

□

A.2 Proofs of results in Sections 4 and 5

Proof of Proposition 1. Let the assumptions of the Proposition be satisfied. (ii) implies (i) is straightforward, so we consider only (i) implies (ii). Since the preference relations are complete on the set of constant acts, they coincide on that set; hence, by the uniqueness clause of the von Neumann-Morgenstern theorem, u_2 is a positive affine transformation of u_2. Hence K and the mapping from $\mathcal{A}$ to $B(K)$ used in the proof of Theorem 1 can be taken to be the same for the two agents; we use the notation employed in that proof. By (i), for every $r \in \mathcal{S}^+$, $\preceq_r^1 \subseteq \preceq_r^2$, and so, by Ghirardato et al. (2004, Proposition A.1), $\mathcal{C}_r^2 \subseteq \mathcal{C}_r^1$. It follows that $\Xi_2 \sqsubseteq \Xi_1$ and $D_2((f,g)) \subseteq D_1((f,g))$ for all $(f,g) \in \mathcal{A}^2$. □

Proof of Proposition 2. The 'if' direction is straightforward. The 'only if' direction is a simple corollary of the proof of Theorem 1. On the one hand, if $\preceq^1$ and $\preceq^2$ are confidence equivalent, they have identical preferences over constant acts (of which they are maximally confident), and hence the same utilities up to positive affine transformation. On the other hand, if they are confidence equivalent, the sets of preferences $\{\preceq_r | r \in \mathcal{S}^+\}$ defined in the proof of Theorem 1 are the same, and so the confidence rankings are the same. □

Proof of Theorem 2. The direction (ii) to (i) is straightforward, so consider the direction (i) to (ii). Part (a) follows from B1 and Theorem 1. As for part (b), consider $f, g \in \mathcal{A}$. If $\min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s) \geqslant \min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s)$, then, for any $c \in \Delta(X)$, if $\sum_{s\in S} u(g(s)).p(s) \geqslant u(c)$ for all $p \in D((f,g))$ then $\sum_{s\in S} u(f(s)).p(s) \geqslant u(c)$ for all $p \in D((f,g))$. Hence, by B3, $f \not<^n g$, whence, by B2, $f \geq^n g$. Similarly, if $\min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s) \leqslant \min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s)$ then $g \geq^n f$. Hence representation (2) holds. Uniqueness follows from Theorem 1.

□

A.3 Proofs of results in Section 6

As stated in Section 6 (see in particular footnote 15), we continue to use the standard notation (and generic terms $f, g \ldots$) for acts, as well as the standard notation (and generic terms $x, z \ldots$) for commodities. In particular, for commodities x, z and $\alpha \in [0,1]$, $\alpha x + (1-\alpha)z$ is the standard vector sum of products of the two commodities, whereas $x_\alpha z$ is the act obtained by applying the mixture operation on (the acts corresponding to) the commodities. Whilst $x_\alpha z$ does not in general belong to $\Re^S_+$, for any preference relation $\preceq^i$ with the properties specified in Section 6, there is a natural element in $\Re^S_+$ corresponding to it; namely $((u^i)^{-1}(\alpha u^i(x_1) + (1-\alpha)u^i(z_1)), \ldots, (u^i)^{-1}(\alpha u^i(x_{|S|}) + (1-\alpha)u^i(z_{|S|})))$. (At each state, the lottery obtained in that state is replaced by its certainty equivalent.) Henceforth we denote this element by $x^i_\alpha z$.

We first require the following Lemma.

Lemma 11. *Under Assumption 1, the strict preferences $\succ^i$ have the following* reduced convexity *property: for all $f, g, h \in \mathcal{A}$, if $g, h \succ^i f$, then, for all $\alpha \in (0,1)$, there exists $\beta \in (0,1]$ such that $(g_\alpha h)_{\beta'} f \succ^i f$ for all $\beta' \in (0,\beta]$.*

Proof. Let $f, g, h \in \mathcal{A}$ such that $g, h \succ^i f$, and consider $\alpha \in (0,1)$. By the monotone decreasing and continuity properties of stakes, there exists $\beta \in (0,1]$ such that $((g_\alpha h)_\beta f, f) \leqq$

$\min\{(g,f),(h,f)\}$. By the representation (1), it follows that $(g_\alpha h)_\beta f \succeq^i f$. However, if $(g_\alpha h)_\beta f \sim^i f$, then, by indifference consistency (and the properties of the representation) $(g_\alpha h)_\gamma h' \sim^i f_\gamma h'$ for any $\gamma \in (0,1]$, $h' \in \mathcal{A}$ with $((g_\alpha h)_\gamma h', f_\gamma h') \equiv \min\{(g,f),(h,f)\}$, contradicting $g, h \succ^i f$, under representation (1). Hence $(g_\alpha h)_\beta f \succ^i f$ and, by indifference consistency, the fact that the stakes are monotone decreasing, and the properties of the representation, $(g_\alpha h)_{\beta'} f \succ^i f$ for all $\beta' \in (0,\beta]$, as required. □

Proof of Theorem 3. For any $x \in \Re^S_+$, let $\pi^i(x) = \{p \in \Delta(\Sigma) | \forall z \in \Re^S_+, \text{if } z \succ x, \text{ then } p \cdot z > p \cdot x\}$, and let $\bar{\pi}^i(x) = \{p \in \Delta(\Sigma) | \forall z \in \Re^S_+, \text{if } z \succ x, \text{ then } p \cdot z \geqslant p \cdot x\}$. On inspection, it is straightforward to check that the reduced convexity property (Lemma 11), combined with the concavity of u and the monotonicity of representation (1), is sufficient for the application of standard arguments on welfare theorems in the absence of completeness and transitivity, notably Fon and Otani (1979), yielding the conclusion that, if x is Pareto optimal, there exists $p \in \bigcap_i \bar{\pi}^i(f^i)$. (In a word, in the presence of reduced convexity and concavity of the utility function, Pareto optimality implies that the convex hull of the strict upper contour set of x^i is disjoint from $\{x^i\}$, allowing application of a separating hyperplane theorem. By monotonicity of representation (1), the separating hyperplane has a positive normal; by normalising, this yields a $p \in \bigcap_i \bar{\pi}^i(x^i)$.) We show that $\pi^i(x) = \bar{\pi}^i(x)$ for all i and $x \in \Re^S_+$. Suppose not, and let $p \in \bar{\pi}^i(x) \backslash \pi^i(x)$ for some i and x; so there exists z with $z \succ^i x$ and $p \cdot z = p \cdot x$. By the fact that stakes are monotone decreasing, indifference consistency of preferences, and representation (1), $z^i_\alpha x \succ^i x$ for any $\alpha \in (0,1]$. By strict concavity of u, for all $s \in S$, $(z^i_\alpha x)_s = (u^i)^{-1}(\alpha u^i(z_s) + (1-\alpha)u^i(x_s)) \leqslant \alpha z_s + (1-\alpha)x_s$, with strict inequality whenever $z_s \neq x_s$. It follows that either $p \cdot x = p \cdot (\alpha z + (1-\alpha)x) > p \cdot (z^i_\alpha x)$, contradicting the assumption that $p \in \bar{\pi}^i(f)$, or $p(s) = 0$ whenever $x_s \neq z_s$. Consider the latter case, and let $S_1 = \{s \in S |\, p(s) = 0\}$. By full support, $\min_{q \in \overline{\bigcup_{\mathcal{C} \in \Xi^i} \mathcal{C}}} \frac{q(S_1)}{q(S \backslash S_1)} > 0$; pick any $\delta > 0$ with $\min_{q \in \overline{\bigcup_{\mathcal{C} \in \Xi^i} \mathcal{C}}} \frac{q(S_1)}{q(S \backslash S_1)} > \delta \frac{\max_{s \notin S_1} u^{i\prime}(z_s)}{\min_{s \in S_1} u^{i\prime}(z_s)}$. For $\epsilon > 0$ and define the allocation z^ϵ as follows: $z^\epsilon_s = \epsilon$ for $s \in S_1$, and $z^\epsilon_s = -\epsilon.\delta$ for $s \notin S_1$. By the definition of z^ϵ, $z + z^\epsilon \succ^i x$ for ϵ sufficiently small, and $p \cdot (z + z^\epsilon) < p \cdot x$ for all $\epsilon > 0$, contradicting the assumption that $p \in \bar{\pi}^i(x)$. Hence $\bar{\pi}^i(x) = \pi^i(x)$ as required.

By standard arguments, if $\bigcap_i \pi^i(x^i) \neq \varnothing$, then $(x^1, \ldots, x^n)$ is Pareto optimal. It remains to show that $\Pi^i(x) = \pi^i(x)$ for all i and $x \in \Re^S_+$.

We first show that $\Pi^i(x) \subseteq \pi^i(x)$. Note that, if $z \succ^i x$, then $\sum_s p(s)(u^i(z_s) - u^i(x_s)) > 0$ for all $p \in ri(D^i((x,z)))$ and hence for all $p \in \bigcap_{z \neq x} ri(D^i(x,z))$. By concavity of u^i, it follows that $\sum_s p(s)u^{i\prime}(x_s)(z_s - x_s) > 0$ for all $p \in \bigcap_{z \neq x} ri(D^i(x,z))$. Renormalising, it follows that, for any $q \in \Pi^i(f)$, $\sum_s q_s \cdot (z_s - x_s) > 0$, and hence that $q \cdot z > q \cdot x$. Since this holds for all $z \in \Re^S_+$ with $z \succ^i x$, $q \in \pi^i(x)$.

We now show that $\pi^i(x) \subseteq \Pi^i(x)$. Suppose not, and let $\bar{p} \in \pi^i(x) \backslash \Pi^i(x)$. Since, as is straightforwardly checked, $\Pi^i(x)$ is convex, by a separation theorem, there exists $y \in \Re^S$ and $b \in \Re$ with $\bar{p} \cdot y \leqslant b \leqslant q \cdot y$ for all $q \in \overline{\Pi^i(x)}$ where the right hand inequality is strict for all $q \in ri(\Pi^i(x))$. Without loss of generality, we can take $b = 0$. Since this implies that, for $\alpha > 0$, $q \cdot \alpha y = \frac{1}{\sum_{t \in S} p(t) u^{i\prime}(x_t)} \sum_s p(s) u^{i\prime}(x_s) \alpha y_s \geqslant 0$ for all $p \in \overline{\bigcap_{z \neq x} ri(D^i(x,z))}$ with strict inequality for all $p \in ri(\bigcap_{z \neq x} ri(D^i(x,z)))$, and since $\overline{\bigcap_{z \neq x} ri(D^i(x,z))}$ is compact, it follows that, for α sufficiently small, $\sum_s p(s)(u^i((x + \alpha y)_s) - u^i(x_s)) \geqslant 0$ for all $p \in \bigcap_{z \neq x} ri(D^i(x,z))$, with strict inequality for some such p. Let A be the set of α possessing this property and such that $x + \alpha y \in \Re^S_+$; we show that $x + \alpha y \succ^i x$ for some $\alpha \in A$. If not, then for every $\alpha \in A$, there exists $\hat{p} \in D^i((x + \alpha y, x))$ with $\sum_s \hat{p}(s)(u^i((x + \alpha y)_s) - u^i(x_s)) < 0$. By nestedness of Ξ^i, it follows that there exists $\hat{p} \in \bigcap_{z \neq x} ri(D^i(x,z))$ with $\sum_s \hat{p}(s)(u^i((x + \alpha y)_s) - u^i(x_s)) < 0$, contradicting the inverse inequality above. Hence $x + \alpha y \succ^i x$ for some $\alpha > 0$, whereas $\bar{p} \cdot (x + \alpha y) \leqslant \bar{p} \cdot x$, so $\bar{p} \notin \pi^i(x)$, as required.

□

Proof of Proposition 3. We show this on the example given in the text. Consider the full insurance allocation $(z^1_{\bar{\delta}}, z^2_{\bar{\delta}}) = ((\bar{\delta} w, \bar{\delta} w), ((1 - \bar{\delta})w, (1 - \bar{\delta})w)$. Agent 2 would accept to exchange x^2 for this ($x^2 \prec^2 z^2_{\bar{\delta}}$) iff:

$$0.5u^2((1 - \delta)w) + 0.5u^2(\delta w) < u^2(1 - \bar{\delta})$$

This gives a strict upper bound ν on $\bar{\delta}$. If a condition analogous to that in (4) holds, namely:

$$\text{(6)} \qquad \min\{\eta w \max\{|\delta - \bar{\delta}|, |\bar{\delta} - (1 - \delta)|\}, 0.45\} + 0.5 > \frac{\bar{\delta}^{1-\gamma^1} - (1 - \delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1 - \delta)^{1-\gamma^1}}$$

for all $\bar{\delta} < \nu$, then, for all $\bar{\delta}$ such that $x^2 \prec^2 z^2_{\bar{\delta}}$, $x^1 \not\prec^1 z^1_{\bar{\delta}}$; hence there are no Pareto optimal allocations accessible from (x^1, x^2). It is straightforwardly checked that, with the parameter values given in the text, these conditions are satisfied.

□

Proof of Proposition 4. It suffices to give an example where no Pareto optimum is m-accessible for any finite m; we use a refinement of the previous example, with $\gamma^1 = \gamma^2 = 1$. Take an allocation $(x^1, x^2) = ((\delta_1 w, \delta_2 w), ((1-\delta_1)w, (1-\delta_2)w))$ with the following properties:

(a) $1 > \delta_1 > \delta_2 > 0$

(b) $\delta_1 - \delta_2 < 18\delta_2(1-\delta_1)$

(c) $\eta w > \max\left\{\frac{1}{(1-\delta_1)+(1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1-1)}, \frac{1}{\delta_1+\delta_1^{0.5}\delta_2^{0.5}-2\delta_1^{0.5}\delta_2^{1.5}}, 2\right\}$.

It is straightforward to see that such allocations exist: $\delta_1 = \frac{3}{4}$, $\delta_2 = \frac{1}{4}$, $\eta = \frac{2.5}{w}$ is an example. Suppose, for reductio, that a Pareto optimal allocation is m-accessible for some finite m: there exists a sequence of allocations $(x_j^1, x_j^2) = ((\delta_{j1}w, \delta_{j2}w), ((1-\delta_{j1})w, (1-\delta_{j2})w))$, $1 \leqslant j \leqslant m+1$, with $(x_1^1, x_1^2) = (x^1, x^2)$, $x_{j+1}^i \succ^i x_j^i$ or $x_{j+1}^i = x_j^i$ for all i, j, and (x_{m+1}^1, x_{m+1}^2) Pareto optimal – and so $(x_{m+1}^1, x_{m+1}^2) = ((\delta' w, \delta' w), ((1-\delta')w, (1-\delta')w))$ for some $\delta' \in [0,1]$. Without loss of generality, it can be assumed that $\delta' \leqslant \delta_{(j+1)1} \leqslant \delta_{j1} \leqslant \delta_1$ and $\delta' \geqslant \delta_{(j+1)2} \geqslant \delta_{j2} \geqslant \delta_2$ for all $1 \leqslant j \leqslant m$. Moreover, such a sequence implies that $0.5u^2((1-\delta_1)w) + 0.5u^2((1-\delta_2)w) < u^2((1-\delta')w)$ and $p(s_1)u^1(\delta_1 w) + p(s_2)u^1(\delta_2 w) < u^1(\delta' w)$ for all $p \in \bigcap_{x' \neq x^1} ri(D^1(x^1, x'))$; since $\bigcap_{x' \neq x^1} ri(D^1(x^1, x'))$ contains (only) the probability measure giving the value 0.5 to each state, it follows that $0.5u^1(\delta_1 w) + 0.5u^1(\delta_2 w) < u^1(\delta' w)$. Hence:

(7) $$\delta_2 < \delta_1^{0.5}\delta_2^{0.5} < \delta' < 1-(1-\delta_1)^{0.5}(1-\delta_2)^{0.5} < \delta_1$$

Consider an arbitrary consecutive pair (x_j^1, x_j^2) and (x_{j+1}^1, x_{j+1}^2) in the sequence. By Theorem 3 (applied in the special case of confidence rankings involving single sets of probability measures) and the fact that the latter is a Pareto-improvement on the former, $\{(\frac{p(s_1)u^{1\prime}(x_j^1(s_1))}{\sum_{t\in S} p(t)u^{1\prime}(x_j^1(t))}, \frac{p(s_2)u^{1\prime}(x_j^1(s_2))}{\sum_{t\in S} p(t)u^{1\prime}(x_j^1(t))} | \; p \in ri(D^1(x_j^1, x_{j+1}^1))\} \cap \{(\frac{p(s_1)u^{2\prime}(x_j^2(s_1))}{\sum_{t\in S} p(t)u^{2\prime}(x_j^2(t))}, \frac{p(s_2)u^{2\prime}(x_j^2(s_2))}{\sum_{t\in S} p(t)u^{2\prime}(x_j^2(t))}t) | \; p \in ri(D^2(x_j^2, x_{j+1}^2))\} = \varnothing$. Doing the calculations, and using the fact that $ri(D^2(x_j^2, x_{j+1}^2)) = 0.5$, this is the case if $\frac{\delta_{j1}-\delta_{j1}\delta_{j2}}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}} \notin ri(D^1(x_j^1, x_{j+1}^1))$. Hence we must have that:

(8) $$\min\{\eta w \max\{|\delta_{j1} - \delta_{(j+1)1}|, |\delta_{(j+1)2} - \delta_{j2}|\}, 0.45\} + 0.5 \leqslant \frac{\delta_{j1} - \delta_{j1}\delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}}$$

Note that $0.95 \leqslant \frac{\delta_{j1}-\delta_{j1}\delta_{j2}}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}}$ if and only if:

$$\begin{aligned}\delta_{j1}-\delta_{j2} &\geqslant 18\delta_{j2}(1-\delta_{j1}) \\ &\geqslant 18\delta_2(1-\delta_1)\end{aligned}$$

by the bounds noted above on δ_{j1} and δ_{j2}. It follows from assumption (b) and the fact that $\delta_{j1}-\delta_{j2} \leqslant \delta_1-\delta_2$ for all j, that, for all j, $0.95 > \frac{\delta_{j1}-\delta_{j1}\delta_{j2}}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}}$. (8) thus reduces to the following inequalities

$$\begin{aligned}\delta_{(j+1)1} &\geqslant \delta_{j1}-\frac{1}{2\eta w}\frac{\delta_{j1}-\delta_{j2}}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}} \\ \delta_{(j+1)2} &\leqslant \delta_{j2}+\frac{1}{2\eta w}\frac{\delta_{j1}-\delta_{j2}}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}}\end{aligned}$$

And so:

$$\delta_{(j+1)1}-\delta_{(j+2)2} \geqslant (\delta_{j1}-\delta_{j2})\left(1-\frac{1}{\eta w}\frac{1}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}}\right) \tag{9}$$

But, using (7):

$$\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2} \geqslant \begin{cases}(1-\delta_1)+(1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1-1) & \text{if } \delta_{j1},\delta_{j2} > \frac{1}{2} \\ \delta_1^{0.5}\delta_2^{0.5}+\delta_2-2\delta_1^{0.5}\delta_2^{1.5} & \text{if } \delta_{j1},\delta_{j2} < \frac{1}{2} \\ \frac{1}{2} & \text{if } (\delta_{j1}-\frac{1}{2})(\delta_{j2}-\frac{1}{2}) \leqslant 0\end{cases}$$

It thus follows from (9) that:

$$\delta_{(j+1)1}-\delta_{(j+1)2} \geqslant (\delta_{j1}-\delta_{j2})(1-\chi) \tag{10}$$

where

$$\chi = \max\left\{\frac{1}{\eta w}\frac{1}{(1-\delta_1)+(1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1-1)}, \frac{1}{\eta w}\frac{1}{\delta_1^{0.5}\delta_2^{0.5}+\delta_2-2\delta_1^{0.5}\delta_2^{1.5}}, \frac{2}{\eta w}\right\}$$

By assumption (c), $\chi < 1$. Iterating inequality (10), we obtain:

$$\delta_{(m+1)1}-\delta_{(m+1)2} \geqslant (\delta_1-\delta_2)(1-\chi)^m > 0$$

contradicting the assumption that (x^1_{m+1}, x^2_{m+1}) is Pareto optimal.

□

B Properties involved in Assumption 1

Readers familiar with the literature on general equilibria in the absence of completeness or transitivity might expect these results to be directly applicable to the case considered in Section 6. This is in fact not straightforwardly possible in general, for two reasons. Firstly, the weakening of transitivity in representation (1) implies that, for certain notions of stakes, preferences represented by (1) may not be convex. (Consider a stakes relation where $(f, g_\alpha h) > (f, g), (f, h)$, for some $\alpha \in (0, 1)$; with such a notion of stakes, the preferences $g, h \succ f \nsucc g_\alpha h$ are compatible with representation (1).) Secondly, since the weak preference order is taken as primitive, it does not follow from the axioms in Section 2.3 that the strict preference order is continuous.[23]

Assumption 1 deals with these issues, by assuming the following properties of the stakes relation and of the preferences.

Monotone decreasing For all $f, g \in \mathcal{A}$ and $\alpha, \beta \in [0, 1]$, if $\alpha \leqslant \beta$, then $(f, g_\alpha f) \leqq (f, g_\beta f)$.

Indifference consistency For all $f, g, h \in \mathcal{A}$ and $\alpha \in (0, 1)$ such that $f \asymp g$ and $f_\alpha h \asymp g_\alpha h$, if $f \sim g$, then $f_\alpha h \sim g_\alpha h$.

Full support For all $s \in S$, there exists $c_s \in \Re_{++}$ such that, for all $h \in \mathcal{A}$ and $\alpha \in (0, 1]$, $(1_s)_\alpha h \succ (c_s)_\alpha h$.[24]

The monotone decreasing property of stakes states that as one considers choices between acts that are "closer" to each other (in the sense of mixtures), the stakes decrease. This property is satisfied by some of the notions of stakes mentioned in Section 2.2 – for example, stakes as the minimum utility which could be obtained from either of the acts, as the difference between the minimum utilities yielded by the two acts, or as the maximum

[23] Schmeidler (1971) has shown that for incomplete transitive preferences (over appropriate spaces), the weak and strict preference orderings cannot both be continuous; although his result does not apply here, due to the weakening of transitivity, it emphasises the subtlety of the issue of the continuity of the derived strict preference ordering. Note that this issue could also have been resolved by taking the strict preference ordering as primitive and using a version of the representation proposed by Bewley (2002); see Section 1.2.

[24] 1_s is the characterisic function for s: $1_s(s') = 1$ for $s' = s$ and $1_s(s') = 0$ for $s' \neq s$. As specified in Section 2.1, c_s is the constant act taking the degenerate lottery yielding c_s for sure in all states.

utility difference between the two acts (taken over the set of states) – whereas it is not satisfied by other notions – for example, taking the maximum utility which could be obtained from either of the acts.

Indifference consistency implies that indifferences cannot become strict preferences simply on altering the stakes. Under representation (1), it is possible that $f \sim g$, because the expected utilities are equal for all $p \in D((f,g))$, but that $f_\alpha h > g_\alpha h$ for some α and h, because $D((f_\alpha h, g_\alpha h)) \supsetneq D((f,g))$ contains a probability measure q such that the expected utility of f calculated with q is greater than that of g. In such cases, $D((f_\alpha h, g_\alpha h))$ is of a higher dimensionality than $D((f,g))$; the latter set, but not the former, is contained in the hyperplane defined by $u \circ f \sim u \circ g$. Indifference consistency rules out such possibilities; technically, it can be thought of as a constant dimensionality assumption on the sets of probability measures in the space $\Delta(\Sigma)$.

Full support is the behavioural formulation of the following full support property of Ξ: for each $s \in S$, there exists $b_s > 0$ such that $p(s) \geqslant b_s$ for all $p \in \bigcup_{\mathcal{C}_i \in \Xi} \mathcal{C}_i$. This property can be thought of as the analogue of full support for a probability measure, but for sets of measures and confidence rankings. In particular, it is stronger than simply asking that all probability measures in the confidence ranking have full support: it requires moreover that probability measures have a common non-zero lower bound on the values for each state.

References

Aliprantis, C. D. and Border, K. C. (2007). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition.

Anscombe, F. J. and Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics*, 34:199–205.

Bewley, T. F. (1989). Market innovation and entrepeneurship: A knightian view. Technical Report 905, Cowles Foundation.

Bewley, T. F. (2002). Knightian decision theory. part i. *Decisions in Economics and Finance*, 25:79–110.

Billot, A., Chateauneuf, A., Gilboa, I., and Tallon, J.-M. (2000). Sharing beliefs: Between agreeing and disagreeing. *Econometrica*, 68(3):685–694.

Dubra, J., Maccheroni, F., and Ok, E. A. (2004). Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133.

Dunford, N. and Schwartz, J. T. (1958). *Linear Operators. Part I: General Theory*. Interscience Publishers, New York.

Faro, J. H. (2009). Variational bewley preferences. Technical report, Cedeplar-FACE-UFMG.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Fon, V. and Otani, Y. (1979). Classical welfare theorems with non-transitive and non-complete preferences,. *Journal of Economic Theory*, 20(3):409–418.

Galaabaatar, T. and Karni, E. (2010). Objective and subjective expected utility with incomplete preferences. Technical report, John Hopkins.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *J. Econ. Theory*, 118(2):133–173.

Gilboa, I., Maccheroni, F., Marinacci, M., and Schmeidler, D. (2010). Objective and subjective rationality in a multiple prior model. *Econometrica*, 78(2).

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2):141–153.

Hill, B. (2009). Confidence in preferences. Technical report, GREGHEC, HEC Paris.

Hill, B. (2010). Confidence and decision. Mimeo, downloadable at https://studies2.hec.fr/jahia/Jahia/site/hec/pid/2563.

Hill, B. (2011). Deferral, incomplete preferences and confidence. Technical report, GREGHEC, HEC Paris.

Kopylov, I. (2009). Choice deferral and ambiguity aversion. *Theoretical Economics*, 4(2):199–225.

Lehrer, E. and Teper, R. (2011). Justifiable preferences. *Journal of Economic Theory*, 146(2):762–774.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74:1447–1498.

Nau, R. (2006). The shape of incomplete preferences. *The Annals of Statistics*, 34(5):2430–2448.

Nau, R. F. (1992). Indeterminate probabilities on finite sets. *The Annals of Statistics*, 20(4):1737–1767.

Nehring, K. (2009). Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory*, 144(3):1054–1091.

Ok, E. A., Ortoleva, P., and Riella, G. (2008). Incomplete preferences under uncertainty: Indecisiveness in beliefs vs. tastes. Technical report, NYU.

Riella, G. and Teper, R. (2011). Probabilistic dominance and status quo biais. Mimeo.

Rigotti, L. and Shannon, C. (2005). Uncertainty and risk in financial markets. *Econometrica*, 73(1):203–243.

Rigotti, L., Shannon, C., and Strzalecki, T. (2008). Subjective beliefs and ex ante trade. *Econometrica*, 76(5):1167–1190.

Schmeidler, D. (1971). A condition for the completeness of partial preference relations. *Econometrica*, 39(2):403–404.

Seidenfeld, T., Schervish, M., and Kadane, J. (1995). A representation of partially ordered preferences. *The Annals of Statistics*, 23(6):2168–2217.