Confidence as a Source of Deferral*

Brian Hill
GREGHEC, HEC Paris & CNRS†

October 10, 2014

Abstract

One apparent reason for deferring a decision – abstaining from choosing, leaving the decision open to be taken by someone else, one's later self, or nature – is for lack of sufficient confidence in the relevant beliefs. This paper develops an axiomatic theory of decision in situations where a costly deferral option is available that captures this source of deferral. Drawing on it, a preliminary behavioural comparison with other accounts of deferral, such as those based on information asymmetry, is undertaken, and a simple multi-factor model of deferral – involving both confidence and information considerations – is formulated. The model suggests that incorporating confidence can account for cases of deferral that traditional accounts have trouble explaining.

Keywords: Confidence; multiple priors; deferral; delegation; information acquisition; value of information; incomplete preferences.

JEL classification: D81, D80, D83.

---

*I would like to thank the public of the "Decision making under uncertainty and beyond" workshop (2013) for useful comments. 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.

†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.

1 Introduction

1.1 Motivation

'If you're not sure, say so.' At first glance, this appears to be a reasonable guide to when a person should *defer*: that is, when he should abstain from choosing any of the options on offer, leaving the decision open, to be taken by someone else, by his future self, or by nature. If you're not sure about any of the options on offer, deferring is one way of saying so. However, this maxim is largely ignored by standard economic accounts of deferral.

Deferral or delegation of decisions to another agent is normally analysed in terms of the expected difference in information (or information-gathering ability) between the 'deferrer' and the 'deferree', without taking into account how sure the former would be in taking the decision himself (for example, Holmstrom (1984)). One family of theories of deferral to one's future self is based on expectations about what information is available to one's future self or what his preferences are (for example, Stigler (1961); Marschak and Miyasawa (1968); Koopmans (1964); Kreps (1979)). Another family considers deciding to be a task, and analyzes deferral – or procrastination of that task – in terms of the comparison between the value of choosing now and the discounted value of choosing at a future moment, in the presence of time inconsistencies or salience effects (for example, Akerlof (May, 1991); O'Donoghue and Rabin (2001)). Under all of these accounts, whether the decision maker is sure or not about which is the best option is assumed not to be a factor in his deferring the decision. The sources of deferral involved are *extrinsic*: they inevitably refer to factors beyond the immediate decision under consideration (the information held by another agent or by one's future self, one's future preferences, the value of accomplishing the task of deciding tomorrow). By contrast, not being sure about what to choose is a source of deferral that only makes reference to the decision maker's own attitudes at the moment of decision: it is *intrinsic*.

Despite their absence from accounts of deferral to others or to one's future self, intrinsic considerations do seem to be involved in some theories of status quo choice – which can be thought of as an instance of deferral to nature. Most prominent are those involving incomplete preferences and according to which the decision maker defers to the status quo in the absence of appropriate determinate preference (for example, Bewley (1986 / 2002)) . To the extent that his preferences are incomplete, he is not sure of the best option, and this plays a role in his deferral to the status quo.

Beyond the plausible intuition that how sure a decision maker is may play a role in his deferral behaviour, the peculiarity of the current state of the literature – where some of our best models of deferral to nature involve intrinsic sources, whilst the main models of deferral to others or to one's future self assume that such factors have no role to play – pleads in favour of the development of theories of intrinsic deferral general enough to cover the latter sorts of cases.

Moreover, intrinsic sources need not be without significant economic consequences, even for deferral to others or to one's future self. Consider, for example, a decision maker faced with an investment decision that can be deferred (or delegated) to a portfolio manager, and suppose that the information differential between the decision maker and the portfolio manager does not justify the manager's fee; mutual fund data suggests that this situation may be quite realistic (Malkiel, 1995; Gruber, 1996). Whilst the information differential and the cost of delegation cannot explain the (observed) deferral in such cases, incorporation of intrinsic sources of deferral can: the decision maker defers because he is not sure how to invest. Or, to take another example, consider a single, important investment decision, among several given options, where all the information one could expect to learn in a reasonable timeframe is already available, and suppose that the decision maker may defer the decision to his future self (or procrastinate); the choice of whether and how to invest in a 401(k) pension plan may be a decision of this sort. Given the absence of a significant information advantage to deferring, and given that, under the time inconsistency-salience approach, there is a lower tendency to procrastinate on single, important tasks – such as deciding one's investment policy for retirement – it may not be straightforward to account for deferral in such cases on the basis of extrinsic sources alone.[1] Once again, incorporation of intrinsic sources of deferral – most people are not very confident in their judgements about the best investment policies – could contribute to understanding deferral of these sorts of decisions. Similarly, deciding on the basis of tealeaves, palm-reading, horoscopes or oracles, or delaying the decision to consult one of these sources, are forms of deferral – to others or to one's future self – for which information acquisition is too scant (even by the lights of many who rely on these methods) and time-consistency considerations too weak

[1] O'Donoghue and Rabin (2001) show that procrastination falls as importance rises when there is a single task to be accomplished. To account for procrastination of important decisions, they develop a model of procrastination in the presence of several tasks. For some further discussion of the relationship to their model, see Section 4.2, footnote 25.

to explain, and where the decision maker's lack of confidence about the appropriate choice may have a role to play.[2] As these examples suggest, in natural situations where decision makers are not sure about what to do – and thus may feel uncomfortable deciding – incorporation of this aspect may help understand cases of deferral that are not clearly accounted for by existing extrinsic approaches.

Any theory of intrinsic deferral with a serious claim of covering all species of deferral – to others, to one's future self and to nature – must be able to tackle the issue of the *cost of deferral*. Deferring to others or to one's future self may incur costs, such as costs of delegation or delay. Standard extrinsic accounts accommodate such costs, roughly speaking, by assigning a value to deferral that is 'weighed off' against its cost. But what is the value of deferral in a situation where the decision maker is not sure what to choose? Existing theories of intrinsic deferral, such as those treating status quo choice, offer no answer to this question.

The aim of this paper is to propose a theory of decision in the presence of a costly deferral option according to which deferral is driven by the decision maker's confidence. This theory will provide the basis for a proper evaluation of this intrinsic source of deferral and its consequences. First of all, since it applies in the costly deferral situations that are the focus of most extrinsic theories of deferral, it will allow a behavioural comparison with these approaches, revealing behavioural differences between them. Secondly, it opens the door to the development of multi-factor models of deferral. Of course, in many real-life situations, several factors – both intrinsic and extrinsic – may be in play. It is standard practice to propose and study models based entirely on one factor (say, information acquisition), ignoring others (for example, time consistency). Likewise, our proposal of a model based solely on intrinsic sources of deferral, ignoring extrinsic sources, is intended as a first step. The ultimate aim is to lay the groundwork for the development of models incorporating both intrinsic and extrinsic sources of deferral. As an illustration, a multi-factor model of deferral, incorporating how sure the decision maker is as well as expected differences in information, will be sketched and some consequences explored.

[2] Note that the 'entertainment value' of consulting an oracle does not explain why people tend to consult them specifically when they are faced with a difficult decision. Casual observation suggests that this is indeed the case: indeed, some divination methods are to be used, we have been told, only when there is an important choice to be made.

1.2 Confidence and the Value of Deferral

As noted previously, the central conceptual difficulty for the development of a theory of intrinsic deferral in the presence of a costly deferral option concerns the question of the value of deferral. To fix ideas, let us restrict ourselves to the framework of decision under uncertainty, where preferences are determined by beliefs and utilities. We focus throughout the paper on the role of the decision maker's confidence in his beliefs, tacitly assuming that the decision maker is fully confident in his utilities. The benchmark in this framework is the model of deferral to nature (or status quo choice) proposed by Bewley (1986 / 2002). It involves a set of probability measures, which can be thought of as representing the beliefs or probability judgements[3] the decision maker is sure of. Moreover, the unanimity decision rule used – 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 – can be thought of as reflecting the maxim with which this paper began: if the decision maker is not sure of the relevant beliefs, then he does not come to a decision. However, as noted above, the model provides no guidance as to how to value the option of deferral.

The solution proposed in this paper starts from the observation that the representation of beliefs by a set of probability measures assumes that being sure is an all-or-nothing affair: either a belief or probability judgement holds for all measures in the set – the decision maker is sure of it – or not. This assumption is unrealistic: one can be more confident of some beliefs than others. As explored in Hill (2013, 2014), the introduction of a graded notion of confidence in beliefs allows the development of more refined theories of decision. In particular, it permits the level of confidence required of a belief for it to play a role in a decision to depend, say, on the importance of the decision. In the case of interest here, such dependence would yield a theory of deferral encompassing the following reasonable maxim: decide when one has sufficient confidence in the relevant beliefs given the importance of the decision, and defer if not. Or, to put it in the terms of the adage with which this paper began: when you're not sure *enough*, say so.

Beyond its role in *driving* deferral, a graded notion of confidence may also have a role to play in *pricing* deferral. For, once introduced, one can talk not only of how confident

[3]By 'probability judgement' we mean a statement concerning probabilities, such as 'the probability of event A is greater than 0.3', held by the agent. We use the terms 'belief' and 'probability judgement' interchangeably in the discussion below.

a decision maker is in certain beliefs, but also of how confident he *would need to be* to decide, given the importance of the decision to be made. It is plausible that, among cases where the decision maker lacks sufficient confidence in the relevant beliefs, the option of deferral is perceived as more valuable in decisions that call for more confidence. As the level of confidence appropriate for the decision goes up, the value of not having to decide – the value of deferring – increases as well. Accordingly, a higher cost of deferral is required to induce the decision maker to decide rather than defer.

Formally, we employ the representation of confidence in beliefs by a *centered confidence ranking* – a nested family of sets of probability measures that includes a singleton set – proposed and defended in Hill (2013). Each set in the family corresponds to a level of confidence, and is understood as representing the beliefs held to that level of confidence. The probability measure in the singleton element of the confidence ranking – called the centre – represents all of the decision maker's beliefs, even those in which he has little confidence. Following the cited paper, the decision maker's attitude to choosing on the basis of limited confidence is represented by a *cautiousness coefficient* – a function assigning to each decision a set in the confidence ranking, interpreted as the appropriate level of confidence (in the decision maker's eyes) for the decision. Finally, we introduce a *cost function* that assigns to each level of confidence a real value, understood as the psychological cost of choosing in a decision requiring that level of confidence when one is not sufficiently confident in the relevant beliefs.

Following the ideas mooted above, we develop a model according to which the decision maker's choice from a menu of Anscombe-Aumann acts A where deferral is possible at a (monetary) cost of x is represented as follows. If there is an optimal element according to the unanimity rule à la Bewley applied with the set of priors picked out by the cautiousness coefficient – that is, if he has sufficient confidence in the relevant beliefs given the decision – then he decides. Otherwise his behaviour depends on the cost of deferral x. If it is outweighed by the cost of deciding – given by the cost function – then he defers. If, on the other hand, the cost of deferral outweighs the cost of deciding, he decides, using all of his beliefs, even those in which he is not very confident (that is, using the probability measure in the singleton element of his confidence ranking).

An axiomatic characterisation of this model is proposed, in which all the aforementioned elements are endogenously derived. The essential axiomatic difference with respect to the standard expected utility model (where there is no deferral) is in appropriate weakenings of the basic choice axioms (*grosso modo*, those playing the role of transitivity, completeness and independence in preference-based models) to allow deferral in some circumstances where standard axioms oblige the decision maker to decide. To that extent, the behavioural properties of the model are quite reasonable. Moreover, the model is behaviourally distinguishable (that is, distinguishable on the basis of the choice patterns it licenses) from a popular theory of costly extrinsic deferral based on the expected information acquisition or asymmetry. Finally, it permits comparative statics analyses of the decision maker's attitudes to deferral, which support the interpretations of the various elements of the model suggested above.

The representation result formally shows how confidence can price deferral, by showing that the confidence-related source of deferral generates a 'cost of deciding', which can be elicited in principle from behaviour. This insight opens the door to the development of multi-factor models of deferral. As an illustration, we formulate and briefly investigate an extension of the model of deferral developed here to incorporate information-related considerations. This model naturally yields a simple additive expression determining deferral, combining the standard value of information and the cost of deciding. It can comfortably accommodate the examples given in the Section 1.1, where important decisions are deferred in the absence of clear extrinsic reasons for deferral; in such cases, the confidence-related cost of deciding will generally be large, and could thus 'tip the balance' in favour of deferral.

The framework and basic notions are introduced in Section 2. The representation is given and axiomatised in Section 3. Section 4 contains a comparison with a standard information-based model of deferral and a brief exploration of a multi-factor model. Related literature not treated elsewhere is discussed in Section 5. Proofs, as well as a comparative statics analysis, are to be found in the Appendices.

2 Preliminaries

2.1 Setup

Throughout the paper, we use a version of the Anscombe and Aumann (1963) framework. Let X, the set of outcomes, be a separable metric space. To simplify the treatment of the cost of deferral, we suppose that outcomes are monetary, and hence take $X = \Re$ (with the standard metric). A *consequence* is a probability measure on X with finite support. $\Delta(X)$ is the set of consequences, endowed with the weak topology. Since the space of Borel probabilities measures over X is metrizable (Billingsley, 2009, p72; Aliprantis and Border, 2007, Theorem 15.12), so is $\Delta(X)$. The operation $\dot{-} : \Delta(X) \times X \to \Delta(X)$ is defined as follows: for all $c \in \Delta(X)$ and $x \in X$, $(c \dot{-} x)(y) = c(y + x)$ for all $y \in X$. $c \dot{-} x$ is the result of 'subtracting' a monetary value x from (each outcome yielded by) a consequence c. A function $u : \Delta(X) \to \Re$ is *zeroed* if $u(0) = 0$ and it is *strictly increasing* if, for all $x, y \in X = \Re$, $x \geqslant y$ if and only if $u(x) \geqslant u(y)$.

Let S be a non-empty finite set of states; subsets of S are called *events*. $\Delta(S)$ is the set of probability measures on S, endowed with the Euclidean topology. The objects of choice are *acts*, defined to be functions from states to consequences. $\mathcal{A}$ is the set of acts, endowed with the inherited product metric. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in [0, 1]$, the mixture $\alpha f + (1 - \alpha)h$ is defined by $(\alpha f + (1-\alpha)h)(s, x) = \alpha f(s, x) + (1-\alpha)h(s, x)$. 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 $\Delta(X)$. Similarly, for any $x \in X$, we shall also use x to denote the constant act yielding the degenerate lottery giving x with probability one.

The decision maker is faced with choices from sets of acts, where deferral is possible but costly.[4] To this end, let $\wp(\mathcal{A})$ be the set of non-empty compact subsets of $\mathcal{A}$, endowed with the Hausdorff metric.[5] We call elements of $\wp(\mathcal{A})$ *menus*. Pointwise mixtures of menus are defined as standard: $A_\alpha h = \{f_\alpha h \mid f \in A\}$, for $A \in \wp(\mathcal{A})$, $h \in \mathcal{A}$, $\alpha \in [0, 1]$. Let $\dagger_x$ be the option of deferring at (non-negative) cost x, and $\mathcal{D} = \{\dagger_x \mid x \in \Re_{\geqslant 0}\}$. An element $(A, \dagger_x) \in \wp(\mathcal{A}) \times \mathcal{D}$ represents the situation in which the decision maker is called upon to choose from menu A where deferral costs x. A *choice correspondence for costly deferral* is a correspondence $\gamma : \wp(\mathcal{A}) \times \mathcal{D} \rightrightarrows \mathcal{A} \cup \mathcal{D}$ – that is, a function $\gamma : \wp(\mathcal{A}) \times \mathcal{D} \to 2^{\mathcal{A} \cup \mathcal{D}} \backslash \varnothing$ – such that $\gamma(A, \dagger_x) \subseteq A \cup \{\dagger_x\}$. The choice correspondence for costly deferral γ delivers for each situation of the sort described a set containing elements in A – which, as standard,

[4]Situations where no deferral option is available can be accommodated in this setup, as limiting cases where the cost of deferral becomes prohibitive.

[5]This metric is defined as follows: for $A, B \in \wp(\mathcal{A})$, $\mathbf{h}(A, B) = \max\{\max_{x \in A} \min_{y \in B} d(x, y), \max_{y \in B} \min_{x \in A} d(x, y)\}$, where d is the metric on $\mathcal{A}$.

are interpreted as those which the decision maker is inclined to choose – or $\dagger_x$ – which represents the decision to defer.

Some further notation shall prove useful. Two menus $A, B \in \wp(\mathcal{A})$ are *extensionally equivalent* if there exists a surjective correspondence[6] $\sigma : A \rightrightarrows B$ such that $\gamma(\{f(s), g(s)\}, \dagger_0) = \{f(s), g(s)\}$ for all $f \in A$, $g \in \sigma(f)$ and $s \in S$; if this holds, we write $A \overset{e.e}{\simeq} B$. This is the natural notion of equivalence for menus if, as is standard in formal theories of decision under uncertainty, one treats two acts as essentially the same whenever they yield consequences between which the decision maker is indifferent in every state.

For a utility function $u : \Delta(X) \to \Re$, a set of probability measures $\mathcal{C} \subseteq \Delta(S)$ and a menu $A \in \wp(\mathcal{A})$, let:

$$(1) \quad \sup(A, u, \mathcal{C}) = \left\{ f \in A \mid \sum_{s \in S} u(f(s)).p(s) \geqslant \sum_{s \in S} u(g(s)).p(s) \;\; \forall p \in \mathcal{C}, \forall\, g \in \mathcal{A} \right\}$$

$\sup(A, u, \mathcal{C})$ is the set of optimal elements of A – those that are ranked better than all other elements in A – according to the unanimity rule with u and $\mathcal{C}$, which ranks an act better than another if the former has higher expected utility for all probability measures in $\mathcal{C}$. Note that if $\mathcal{C}$ is a singleton, then $\sup(A, u, \mathcal{C})$ is the set of acts in A with maximal expected utility calculated with u and the element in $\mathcal{C}$.

As a final piece of notation, let $\Phi = \{A \in \wp(\mathcal{A}) \mid \exists h \in \mathcal{A},\ \alpha \in (0, 1], \dagger_0 \in \gamma(A_\alpha h, \dagger_0)\}$. This is the set of menus such that, for at least one mixture of the menu, the decision maker defers.

2.2 *Confidence ranking and cautiousness coefficient*

We adopt with some modification two notions that were introduced in Hill (2013).

Definition 1. A *confidence ranking* Ξ is a nested family of closed, convex subsets of $\Delta(S)$. A confidence ranking Ξ is *continuous* if, for every $\mathcal{C} \in \Xi$, $\mathcal{C} = \overline{\bigcup_{\Xi \ni \mathcal{C}' \subsetneq \mathcal{C}} \mathcal{C}'} = \bigcap_{\Xi \ni \mathcal{C}' \supsetneq \mathcal{C}} \mathcal{C}'$. Ξ is *strict* if, for every $\mathcal{C}_1, \mathcal{C}_2 \in \Xi$ with $\mathcal{C}_1 \subset \mathcal{C}_2$, $(\mathcal{C}_1 \cap (ri(\mathcal{C}_2))^c) \cap ri(\overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}) = \varnothing$.[7] Ξ is *centered* if it contains a singleton set; in this case, the member of the singleton set is called the *centre* and is denoted p_Ξ.

[6]A correspondence $\sigma : A \rightrightarrows B$ is surjective if, for all $g \in B$, there exists $f \in A$ with $g \in \sigma(f)$.

[7]For a set X, $\overline{X}$ is the closure of X, and $ri(X)$ is its relative interior.

As mentioned in the Introduction, confidence rankings represent decision makers' beliefs and their confidence in their beliefs. A set in the confidence ranking corresponds to a level of confidence and can be thought of as representing the beliefs held to this level of confidence. For further discussion of the main features of the definition, the reader is referred to Hill (2013). The only new property in the definition is strictness, which implies that the confidence ranking is strictly increasing in a particular sense. It ensures, for example, that as the confidence level increases, the highest value for the probability of an event that is endorsed at that level of confidence also increases, unless this value is already maximal over all confidence levels.

As noted in Hill (2013), 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, although he may not be very confident in it. He is, so to speak, a 'Bayesian with confidence'. Whilst we by no means wish to suggest that all decision makers are of this sort, we focus on decision makers with centred confidence rankings to facilitate the comparison with other approaches to deferral, which generally assume expected utility.[8]

The second notion required is that of a *cautiousness coefficient* for a confidence ranking Ξ, which is defined to be a function $D : \wp(\mathcal{A}) \to \Xi$ satisfying the following three properties.

(Extensionality) For all $A, B, \in \wp(\mathcal{A})$, if $A \overset{e.e.}{\simeq} B$, then $D(A) = D(B)$.

(Continuity) For all $\mathcal{C} \in \Xi$, the sets $\{A \in \wp(\mathcal{A}) | \; D(A) \supseteq \mathcal{C}\}$ and $\{A \in \wp(\mathcal{A}) | \; D(A) \subseteq \mathcal{C}\}$ are closed.

(Φ-Richness) For all $A \in \Phi \subseteq \wp(\mathcal{A})$ and $\mathcal{C} \in \Xi$, there exists $h, h' \in \mathcal{A}$ and $\alpha, \alpha' \in (0, 1]$ such that $D(A_\alpha h) \subseteq \mathcal{C} \subseteq D(A_{\alpha'} h')$.

The cautiousness coefficient can be understood as assigning a level of confidence to a menu: $D(A)$ represents the beliefs held to the level of confidence appropriate for use in the choice from the menu A. As discussed in Hill (2013, 2014), it is a subjective element that captures the decision maker's attitude to choosing on the basis of limited confidence.[9] The

[8]Possible extensions, involving weakening of the centering and strictness properties in particular, are discussed in Remark 2 (Section 3.1).

[9]Where the cited papers speak of 'attitude to choosing in the absence of confidence', we prefer 'attitude to choosing on the basis of limited confidence' which strikes us as more in tune with the graded notion of confidence involved.

underlying idea is that the appropriate level of confidence is picked out on the basis of the importance of the decision: for more important decisions, more confidence is required. By contrast with Hill (2013, 2014), who exogenously assume a notion of stakes or importance of a decision, this is left implicit in the notion of cautiousness coefficient used here.[10] Accordingly, several new properties of the cautiousness coefficient are required.

Extensionality states that all that counts in the determination of the level of confidence appropriate for a decision are the values of the consequences of the acts in the menu at the different states. Virtually all formal theories of decision under uncertainty treat extensionally equivalent acts – those that yield consequences between which the decision maker is indifferent in every state – as being essentially the same. Extensionality says that whenever two menus are composed of such acts, the same level of confidence is appropriate for the choice from the menus. Continuity, which is fairly standard, seems reasonable: the level of confidence appropriate for choice from a menu may be altered as the menu changes, but one would not expect it to 'jump' with gradual modifications of the menu.

Φ-Richness is a technical property, which states that the appropriate level of confidence for the choice from a menu can be shifted as far up or down as desired, by considering suitable mixtures of the menu. There is a sense in which (in particular in the presence of an independence axiom; see Section 3.2) the choice from A and the choice from $A_\alpha h$ are the 'same' choice. Nevertheless, these choices need not be of the same importance; accordingly, different levels of confidence may be appropriate for the two choices. To that extent, the latter choice can be thought of as a 'version' of the former choice which calls for the use of beliefs held to the level of confidence appropriate for $A_\alpha h$ rather than A. Φ-Richness simply states that for any choice for which the decision maker defers for some version of the choice (that is, every menu in Φ) and any confidence level, there is a version of the choice, obtained by mixing with an act, for which the appropriate confidence level is above the level in question, and there is a version for which the appropriate confidence level is below that level. The intuition is that mixing with an act can change many of the properties of a menu, and in particular the main properties that are relevant for the importance of the choice from it, and for the level of confidence appropriate.

Remark 1. In Hill (2013, 2014), the notion of the stakes involved in a decision was fixed

[10]A notion of stakes can however be defined from the notion of cautiousness coefficient used here; see Remark 1 below for further discussion.

exogenously: in the former paper by stipulating a notion of stakes, and in the latter paper by assuming an exogenously-given stakes relation which orders decisions according to whether the stakes involved in them are higher or lower. The cautiousness coefficient was defined in these papers as a function that respects the stakes relation. The model developed below can be formulated in terms of a stakes relation and a cautiousness coefficient that respects it (in the style of Hill (2014)): it suffices to take D as given and define the stakes relation $\leqq$ on menus by $A \leqq B$ iff $D(A) \subseteq D(B)$. (The stakes relation defined in this way will automatically satisfy versions of the properties given in Hill (2014).) To that extent, the notion of stakes assumed in the cited papers can be thought of as endogenised, or elicited from choice, in the main theorem of this paper (Theorem 1). For further discussion, see Remark 3 (Section 3.3) and Section 5.

3 A model of costly deferral

We now introduce a representation of choice in the presence of a costly deferral option that captures the role of confidence. The aim in this section is not to offer an exhaustive study of deferral, but rather to propose a relatively simple model of intrinsic deferral that can accommodate cases where deferral is costly, and characterise its behavioural properties.

3.1 Representation

Consider the following representation of a choice correspondence for costly deferral γ: for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$,

$$(2) \qquad \gamma(A, \dagger_x) = \begin{cases} \sup(A, u, D(A)) & \text{if } \sup(A, u, D(A)) \neq \varnothing \\ \sup(A, u, \{p_\Xi\}) & \text{if } \sup(A, u, D(A)) = \varnothing \text{ and } c(D(A)) \leqslant u(x) \\ \{\dagger_x\} & \text{otherwise} \end{cases}$$

where $u : \Delta(X) \to \Re$ is a von Neumann-Morgenstern utility function, Ξ is a continuous strict, centred confidence ranking, $D : \wp(\mathcal{A}) \to \Xi$ is a cautiousness coefficient and $c : \Xi \to \Re_{\geqslant 0} \cup \{\infty\}$ is a continuous function that is order-preserving and -reflecting with respect to $\subseteq$: that is, $c(\mathcal{C}) \geqslant c(\mathcal{C}')$ if and only if $\mathcal{C} \supseteq \mathcal{C}'$.[11] The function c is called the *cost*

[11]We take the standard topology on the non-negative extended reals $\Re_{\geqslant 0} \cup \{\infty\}$, namely that produced by the Alexandroff one-point compactification of $\Re_{\geqslant 0}$ (Aliprantis and Border, 2007, Section 2.16).

function. It can be understood as assigning to any confidence level the psychological cost of deciding – or equivalently the value of deferring – in a choice that calls for that level of confidence but in which one does not hold the relevant beliefs with sufficient confidence to yield a decision. As one would expect, this function is increasing in the confidence level: for decisions requiring more confidence, it is more (psychologically) costly to decide when one would have wanted to defer for lack of sufficient confidence in the appropriate beliefs.

To understand representation (2), first note that, when deferral is costless, it becomes:

$$\gamma(A, \dagger_0) = \begin{cases} \sup(A, u, D(A)) & \text{if } \sup(A, u, D(A)) \neq \varnothing \\ \{\dagger_0\} & \text{otherwise} \end{cases}$$

Recall that $D(A)$ represents the set of beliefs that the decision maker holds to the level of confidence appropriate for use in the choice from the menu A. $\sup(A, u, D(A))$ contains the acts in A that are better than all other alternatives under the unanimity rule with the set of probability measures $D(A)$. It can be interpreted as containing those acts that the decision maker can conclude to be better than all other acts on the menu, on the basis of the beliefs in which he has sufficient confidence given the choice to be made. Hence, under this representation, the decision maker decides when there is an act that is optimal according to these beliefs. If there is no such act, he defers. When deferral is free, a decision maker represented by (2) chooses when he has sufficient confidence in the appropriate beliefs given the decision to be made, and defers when he does not.

In the light of this, representation (2), applied in the general case of choice from a menu A where it costs x to defer, can be understood as follows. If the decision maker has sufficient confidence, his behaviour is the same as in the case where deferral is free: he chooses an optimal act according to the unanimity rule with appropriate beliefs. In the absence of sufficient confidence – when there is no optimal act according to the rule – his behaviour may differ. In these cases, he compares the psychological cost of deciding given the confidence level appropriate for the decision (given, in utility units, by $c(D(A))$) with the cost of deferral in the choice situation in which he is in (which corresponds to the utility of x, $u(x)$). If the monetary cost of deferral imposed on him is lower than the psychological cost of deciding, then he defers, as recommended by the unanimity rule applied with the appropriate confidence level. If the monetary cost of deferral imposed on him outweighs the psychological cost of deciding, then he decides. In these cases, he chooses an act which has highest expected utility calculated with the probability measure in the centre of the confidence ranking. To the extent that this probability measure captures all of the decision maker's beliefs, this amounts to deciding on the basis of all of his beliefs, irrespective of the confidence with which they are held. The intuition here is that, whilst the decision maker should not decide on the basis of a belief in which he has insufficient confidence when he can defer, if deferral is too costly he may as well mobilise all of his beliefs – even those held with little confidence. Given that he is relying on more beliefs, the decision maker may be able to come to a decision; in fact, for decision makers with centered confidence rankings – which, recall, are basically Bayesians with confidence – this will always be the case.

Whilst not the only way of choosing in such situations, the procedure formalised in representation (2) is certainly not unreasonable. The decision maker will choose whenever he has sufficient confidence in the relevant beliefs. Otherwise, he will allow himself to defer only if the (psychological) cost of deciding outweighs the (monetary) cost of deferral. Finally, when he does not allow himself to defer, he decides on the basis of all of his beliefs, even those in which he has little confidence.

Remark 2. Although we consider a simple rule for choice in the face of costly deferral, it is possible to formulate models, and extend the representation results, to incorporate modifications or refinements.

First of all, one could abandon the assumption that the confidence ranking is centered. In the case where the decision maker has insufficient confidence in the relevant beliefs but decides due to the cost of deferral, one could adopt the same strategy of employing the smallest set of probability measures in the confidence ranking. Since this may be a non-singleton set, some other rule, such as the maxmin EU rule (Gilboa and Schmeidler, 1989), would be required in cases in which the unanimity rule applied on this set does not yield an optimal element.

Secondly, one could abandon the use of the smallest set in the confidence ranking to choose in cases where the decision maker has insufficient confidence but the cost of deferral is high. One possibility is to use the set of probability measures corresponding to the level of confidence appropriate for the choice from the menu, and a decision rule that always yields an optimal act, such as the maxmin EU rule. Another possibility is to use the largest set in the confidence ranking that yields a decision (that is, for which the optimal set of acts under the unanimity rule is non-empty). Whenever the confidence ranking is strict, this latter possibility is equivalent to representation (2).

Working in a preference framework, Hill (2014) distinguishes and characterises axiomatically ways of deciding based on all of one's beliefs and ways that rely on the beliefs that one holds to the level of confidence appropriate for the decision. Similar techniques could be employed to generalize Theorem 1 to non-centred or non-strict confidence rankings, and to alternative strategies for choosing in the absence of sufficient confidence.

Thirdly, one could imagine using a criterion for deferral different from the one involved in representation (2), namely $c(D(A)) \leqslant u(x)$. Note that representations with criteria such as $c(D(A)) \leqslant x$ or $c(D(A)) \leqslant -u(-x)$ (where, in the former, c gives the cost in monetary units) are related to representation (2) by appropriate transformations of the cost function, and so are behaviourally indistinguishable from it. Hence the results below immediately apply to them. Other possible suggestions include taking the difference in cost between the confidence level required for the decision and the highest confidence level at which the decision maker holds sufficient beliefs to decide, or the following criterion based on the difference in the value of deciding after paying the monetary cost of deferral and deciding in the absence of this cost: $c(D(A)) \leqslant \max_{h\in A}\sum_{s\in S} u(h(s)).p_{\Xi}(s) - \max_{h\in A}\sum_{s\in S} u(h(s) \dot{-} x).p_{\Xi}(s)$. Representation theorems for versions of representation (2) involving these criteria can be developed, using similar techniques to those adopted in the proof of Theorem 1. Finally, one could go beyond monetary outcomes, developing similar representations and results for outcome spaces and costs that are not purely monetary (including, for example, a temporal element).

3.2 Axioms

To state the axioms, the following definition shall prove useful.

Definition 2. For each $x \in \Re_{\geqslant 0}$, the function $\bar{\gamma}^x : \wp(\mathcal{A}) \to 2^{\mathcal{A}}$ is defined by: $f \in \bar{\gamma}^x(A)$ if and only if, for all $h \in \mathcal{A}$ and $\alpha \in (0,1]$, if $\dagger_x \notin \gamma(A_\alpha h, \dagger_x)$, then $f_\alpha h \in \gamma(A_\alpha h, \dagger_0)$.

To explain the underlying intuition, let us say that the cost x *is not motivating* for a menu A when, for every version of the menu,[12] if the decision maker chooses from it when deferral costs x, then he also chooses when deferral is free. A cost that is not motivating

[12] Recall from the discussion in Section 2.2 that different versions of a menu are obtained by mixing the menu with an act.

for a menu does not drive any decision taken from it: whenever the decision maker decides at this cost of deferral, he is 'sure enough' to decide even when deferral is costless. By contrast, a cost that is motivating for a menu may drive the decision from it: the decision maker may choose at this cost though he does not choose at some lower cost. So $f \in \bar{\gamma}^x(A)$ says that the cost x is not motivating for A and (the appropriate mixture of) f is among the potential choices from (the corresponding version of) A whenever a decision is made. Note that $\bar{\gamma}^x$ may take as value the empty set; it does so on menus for which x is motivating. Hence $\bar{\gamma}^x$ is not a choice correspondence (which is standardly assumed to take non-empty values).

We now consider several axioms on γ, which are organised into three groups.

3.2.1 Main Behavioral Axioms

Axiom A1 (Contraction)**.** For all $A, B \in \wp(\mathcal{A})$ with $A \subseteq B$, all $x \in \Re_{\geqslant 0}$, and all $f \in \mathcal{A}$, if $f \in \bar{\gamma}^x(B)$, then $f \in \bar{\gamma}^x(A)$.

Axiom A2 (Strong Expansion)**.** For all $A, B \in \wp(\mathcal{A})$, $x \in \Re_{\geqslant 0}$, $f \in A$ and $g \in A \cap B$, if $f \in \bar{\gamma}^x(A)$ and $g \in \bar{\gamma}^x(B)$, then $f \in \bar{\gamma}^x(A \cup B)$.

Axiom A3 (Independence)**.** For all $A \in \wp(\mathcal{A})$, $h \in \mathcal{A}$ and $\alpha \in (0,1]$ and all $x, x' \in \Re_{\geqslant 0}$ such that $\dagger_x \notin \gamma(A, \dagger_x)$ and $\dagger_{x'} \notin \gamma(A_\alpha h, \dagger_{x'})$, $\gamma(A, \dagger_x)_\alpha h = \gamma(A_\alpha h, \dagger_{x'})$.

Axiom A4 (Consistency)**.** For all $A \in \wp(\mathcal{A})$, and all $x, x' \in \Re_{\geqslant 0}$ with $x \leqslant x'$, if $\dagger_x \notin \gamma(A, \dagger_x)$ then $\gamma(A, \dagger_{x'}) = \gamma(A, \dagger_x)$.

Axiom A5 (Centering)**.** For all $A \in \wp(\mathcal{A})$, there exists $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(A_\alpha h, \dagger_0)$.

Contraction (A1) is just Sen's axiom α applied to $\bar{\gamma}^x$; as such, it is well known in the literature as a standard ingredient in the revealed preference theory of complete preferences. Similarly, Strong Expansion (A2) is the property π introduced by Hill (2012), formulated for $\bar{\gamma}^x$. As discussed in Hill (2012), π can be thought of as the equivalent of Sen's axiom β for incomplete preferences. Firstly, the intuition supporting it is similar to that supporting the standard axiom β. Secondly, it implies β, and is in fact equivalent to it on ordinary choice correspondences. Finally, α and π basically characterise the rationalisation of a generalisation of standard choice correspondences by reflexive, transitive but not necessarily complete binary relations. Lest it help the reader get a grasp on these axioms, one can think of α and π as boiling down, in the case of binary menus (and hence standard preferences), to the assumption of reflexivity and transitivity without completeness.

To appreciate the formulation of these conditions for $\bar{\gamma}^x$, it is useful to consider the application of Strong Expansion (A2) to menus $\{f, g\}$, $\{g, h\}$ and $\{f, g, h\}$.[13] It implies in particular that, if f is chosen over g, g is chosen over h, and the cost x is not motivating for these decisions – that is, whenever the decision maker chooses at cost x, he also chooses when deferral is free – then this cost cannot be motivating for the choice from $\{f, g, h\}$ – if he chooses at cost x, then he must also choose when deferral is free. The underlying intuition is reasonable: if the decisions taken between f and g and between g and h at cost of deferral x are driven not by the cost of deferral, but say by the decision maker being 'sure enough', then the same holds for the choice from $\{f, g, h\}$.

By contrast, consider the following formulation of π for choice behaviour when deferral is free, which is arguably a more standard version of the property in the current framework:

Axiom A2' (Strong Expansion$_{Free}$)**.** For all $A, B \in \wp(\mathcal{A})$, $f \in A$ and $g \in A \cap B$, if $f \in \gamma(A, \dagger_0)$ and $g \in \gamma(B, \dagger_0)$, then $f \in \gamma(A \cup B, \dagger_0)$.

Applied to the case discussed above, Strong Expansion$_{Free}$ (A2') implies that, if the decision maker chooses f over g and g over h when deferral is free, then he cannot defer from $\{f, g, h\}$ when deferral is free. This may be unreasonably strong in some circumstances. For example, suppose that g represents taking out \$100 of credit to buy a financial product, where the current portfolio is given by h, and f itself represents buying another \$100 product on credit with respect to current portfolio g. Whilst one may be able to decide in each of the binary choices between f and g and g and h, when the credit under consideration is larger (up to \$200 over one's current portfolio) the decision may be more difficult, and it might not be unreasonable to defer.[14] Strong Expansion (A2) can comfortably accommodate such behavioural patterns: it allows deferral in such cases, albeit under certain conditions. In particular, it only prohibits deferral if the decision maker chooses

[13]This case has the advantage of highlighting the relation to the transitivity axiom on preferences. Similar considerations apply to Contraction (A1).

[14]Consider also the version of this example that involves the successive application of Strong Expansion$_{Free}$ to 100 binary choices in which one decides between taking \$100 credit over one's current portfolio or not.

from $\{f, g, h\}$ at a cost of deferral that is not motivating for the other choices (because A2 implies that this cost is not motivating for the choice from $\{f, g, h\}$ either).

Note moreover that, whenever a choice is made from $\{f, g, h\}$, then A2 agrees with A2' that f will be among the chosen options. The only 'violations' of the more standard Strong Expansion$_{Free}$ that are sanctioned by Strong Expansion are not forms of inconsistent choices (such as those corresponding to preference cycles), but cases of deferral where the standard axiom would have demanded decision. Summing up in terms of preferences for readers more familiar with these: on binary menus, Contraction and Strong Expansion basically amount to the assumption of reflexivity and a weakening of transitivity to allow indeterminacy in some cases where the standard axiom demands determinate preference.

Independence (A3) demands that the standard independence axiom, formulated in a choice-theoretic setting, holds whenever the decision maker decides rather than deferring. Evidently it fully retains the intuitions behind the standard axiom, and is equivalent to it in the case where the decision maker never defers. The restriction to pairs of menus from which the decision maker decides is a central behavioural difference between the proposed model and more standard ones. Indeed, in the presence of the other axioms, dropping the condition that $\dagger_{x'} \notin \gamma(A_\alpha h, \dagger_{x'})$ from A3 yields the standard expected utility representation (formulated in the current framework), under which the decision maker never defers. In the current context, this condition is entirely natural: without it, the axiom would demand that if one chooses from A when deferral costs \$100 then one cannot defer from $A_\alpha h$ when deferral is free.

Consistency (A4), requires an immediately intuitive relationship between decisions with different costs of deferral. It says that, if one is willing to choose from a menu at a particular cost of deferral, then as deferral becomes more expensive, one will continue to choose, and make the same choice. Centering (A5) characterises the centeredness property of the confidence ranking (Section 2.2). Recall from the discussion in Section 2.2 that different versions of a choice can be obtained by mixing the menu with an act; these different versions may be of varying importance, with different levels of confidence appropriate. In the light of this, Centering states that, for any menu, there is a version of the choice from the menu in which the decision maker will decide rather than deferring. The basic intuition is that, if the appropriate level of confidence is sufficiently low, the decision maker will hazard a choice, albeit one in which he may not be very confident. As mentioned in Section 2.2, the centeredness property of confidence rankings is adopted to ease comparison with standard approaches to deferral; the same goes for the Centering axiom, which will be dropped in any extension of the model that does not involve centered confidence rankings.

3.2.2 Remaining Behavioral Axioms

Axiom A6 (Defer or Choose)**.** For all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$, if $\dagger_x \in \gamma(A, \dagger_x)$, then $\gamma(A, \dagger_x) = \{\dagger_x\}$.

Axiom A7 (No Deferral)**.** For all $A \in \wp(\mathcal{A})$ with $A \subseteq \Delta(X)$ and all $x \in \Re_{\geqslant 0}$, $\dagger_x \notin \gamma(A, \dagger_x)$.

Axiom A8 (Monotonicity)**.** For all $f, g \in \mathcal{A}$, $z, z' \in X$, $A, B \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$:

i. if $z > z'$ then $\gamma(\{z, z'\}, \dagger_0) = \{z\}$;

ii. if $g(s) \in \gamma(\{f(s), g(s)\}, \dagger_0)$ for all $s \in S$, then $g \in \gamma(\{f, g\}, \dagger_0)$;

iii. if $A \overset{e.e}{\simeq} B$, then $\dagger_x \in \gamma(A, \dagger_x)$ if and only if $\dagger_x \in \gamma(B, \dagger_x)$.

Defer or Choose (A6), which is specific to the sort of choice situation under consideration here, basically says that if the decision maker defers – $\dagger_x$ is in the choice set – then he does not choose – no acts are retained as admissible choices. No Deferral (A7) simply states that in choices among constant acts only, the decision maker never defers. The axiom 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.

The first clause of Monotonicity (A8) is a standard assumption in the context of monetary outcomes, saying that the decision maker chooses (strictly) more money over less. It is retained to simplify the representation (insofar as it implies that different costs of deferral are perceived as different by the decision maker), and can be dropped without significant changes to the basic form of the result below. The second clause is a choice-theoretic version of the standard monotonicity axiom for decision under uncertainty as formulated in the preference framework. The final clause demands that whether the decision maker defers at a given cost of deferral respects extensional equivalence of menus: for any two menus such that each act in one is extensionally equivalent to an act in the other (they yield consequences between which the decision maker is indifferent in all states), the decision maker defers from one menu at a given cost of deferral precisely when he defers from the other menu at this cost. Virtually all formal theories of decision under uncertainty treat extensionally equivalent acts in the same way; this axiom extends this to the case of deferral. (The extension to the case of deferral is not a consequence of clause A8 part ii., as is standardly the case, because of the weakness of the choice-theoretic axioms.)

3.2.3 Technical Axioms

Before stating the final, technical, axioms, let us introduce the following terminology. Define $\iota : \wp(\mathcal{A}) \to \Re_{\geqslant 0} \cup \{\infty\}$ by $\iota(A) = \inf\{x \in \Re_{\geqslant 0} \mid \dagger_x \notin \gamma(A, \dagger_x)\}$,[15] and $mmc : \wp(\mathcal{A}) \to \Re_{\geqslant 0} \cup \{\infty\}$ by $mmc(A) = \inf\{x \in \Re_{\geqslant 0} |\ \bar{\gamma}^x(A) = \varnothing\}$. Moreover, let $\Upsilon = \{A \in \wp(\mathcal{A}) \mid \dagger_0 \in \gamma(A, \dagger_0)\}$.

Υ is the set of menus from which the decision maker defers when deferral is free. For such menus, the cost of deferral counts: whether the decision maker chooses or defers depends on it. ι gives the point where the cost of deferral 'bites': whenever the cost is below the value given by ι the decision maker defers, whereas above that value he decides. Naturally, if a cost is above the 'biting point' for such a menu, then it is motivating for the menu, in the sense introduced previously. The interpretation of ι as a biting point only holds on Υ (ι takes the value zero elsewhere), but mmc can be thought of as an extension beyond this set. It gives the lowest cost that can be motivating,[16] and hence the lowest biting point over all versions of the menu where the cost counts. For any menu A whose biting point is above $mmc(A)$, this point is given by $\iota(A)$, and A is in Υ. So for a menu A not in Υ, although the biting point is not pinned down by $\iota(A)$, one can nevertheless conclude that it is not greater than $mmc(A)$. Note that Φ (defined in Section 2.1) is the set of menus for which some cost is motivating.

Now consider the following axioms.

Axiom A9 (Continuity)**.** For all $A, A_n \in \wp(\mathcal{A})$, $f, f_n, g, h \in \mathcal{A}$, $\alpha \in (0,1)$, $x, x_n \in \Re_{\geqslant 0}$:

i. if $A_n \to A$, $f_n \to f$ and $x_n \to x$, and $f_n \in \gamma(A_n, \dagger_{x_n})$ for all $n \in \mathbb{N}$, then $f \in \gamma(A, \dagger_x)$;

[15]Throughout, we take an infimum to be infinite when the set over which it is taken is empty.

[16]Hence the terminology: mmc stands (albeit loosely) for minimal motivating cost. Recall from the remarks following Definition 2 that $\bar{\gamma}^x(A) = \varnothing$ when x is motivating for A.

ii. if $\sup \iota(\Upsilon) > mmc(\{f, g\})$, $f \in \gamma(\{f, g\}, \dagger_0)$, $f(s) \in \gamma(\{f(s), h(s)\}, \dagger_0)$ and $g(s) \in \gamma(\{g(s), h(s)\}, \dagger_0)$ for all $s \in S$, then there exists $y \in \Re_{\geqslant 0}$ such that $\bar{\gamma}^y(\{f, g\}) = \varnothing$ but $\bar{\gamma}^y(\{f, g_\alpha h\}) \neq \varnothing$;

iii. if $x \in (0, \sup \iota(\Upsilon))$, $(\iota^{-1}([0, x]) \cap \Upsilon) \cup (mmc^{-1}([0, x]) \cap \Upsilon^c)$ and $\iota^{-1}([x, \infty]) \cup (mmc^{-1}([x, \infty]) \cap \Upsilon^c)$ are closed in Φ;

iv. $\{x \in \Re_{\geqslant 0} \mid \bar{\gamma}^x(A) = \varnothing\} = \{x \in \Re_{\geqslant 0} \mid \bar{\gamma}^x(A_\alpha h) = \varnothing\}$ and the set is open in $\Re_{\geqslant 0}$.

Axiom A10 (Richness)**.** For all $A \in \Phi$:

i. for all $x \in \iota(\Upsilon)$, there exist $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ such that $\dagger_y \in \gamma(A_\alpha h, \dagger_y)$ for all $y < x$;

ii. if there exists $x > \inf \iota(\Upsilon)$ with $\bar{\gamma}^x(A) \neq \varnothing$, then $\{A_\alpha h \in \Phi \mid \alpha \in (0, 1], h \in \mathcal{A}\} \nsubseteq \overline{\Upsilon}$.

As indicated, these are basically technical axioms. The first clause of Continuity (A9) is a version of the standard upper hemi-continuity property for correspondences, with the addition of continuity under changes in the cost of deferral. It captures essentially the same intuition that small changes in the choice situation – in the menu and the cost of deferral – do not induce large changes in choice behaviour. The second clause (A9 part ii.) states that, for pairs of acts f and g where f is chosen over g when deferral is free, whenever g is 'worsened' – by mixing it with an act dominated statewise by g and f – some costs cease to be motivating. As g is 'worsened' the choice becomes 'easier', so some costs that could drive the initial decision – in the sense that the decision maker is not sure enough to decide when deferral is free – no longer drive the 'easier' decision – he is sure enough in that decision to decide when deferral is free. The third clause (A9 part iii.) requires a certain continuity in the point at which the cost of deferral bites: as one gradually changes the menu, its biting point (as given by the value of ι if the cost of deferral counts, and the lowest motivating cost mmc if not) does not suddenly 'jump'. The final clause (A9 part iv.) demands that the set of costs that are motivating for a menu is the same for the menu and any version of it, and that this set is open. The assumption of openness, like some of the other parts of the technical axioms, is largely conventional: appropriate modifications of it would lead, for example, to a representation with the same general form as (2) but with a strict rather than weak inequality in the condition on the second line.

Richness (A10) is also a technical axiom. For any menu for which some cost is motivating, the first clause (A10 part i.) demands that the point at which cost of deferral bites can be moved as high up the range of possible values as desired by taking the appropriate version of the menu. The second clause (A10 part ii.) states that either every cost is motivating for the menu – and hence the biting point can be moved as low down as desired by considering different versions of the menu – or a condition is satisfied. The condition basically states that for some versions of the menu, the biting point is not entirely fixed by deferral behaviour: the cost does not count for the choice from them (and they are not arbitrary close to menus for which the cost counts), so ι does not give the biting point. In other words, A10 part ii. says that, if it cannot be concluded that the biting point can be moved as low down the range of possible values as desired (there is a cost that is not motivating for the menu), then it cannot be concluded that it *cannot* be moved as low down the range of values as desired (for some versions of the menu, the cost does not count, so the biting point cannot be fully pinned down).

3.3 Result

We have the following representation result.

Theorem 1. *Let γ be a choice correspondence for costly deferral on $\mathcal{A}$. The following are equivalent:*

(i) *γ satisfies A1–A10,*

(ii) *there exists a strictly increasing zeroed continuous affine utility function $u : \Delta(X) \rightarrow \Re$, a continuous strict centred confidence ranking Ξ with centre p_Ξ, a cautiousness coefficient $D : \wp(\mathcal{A}) \rightarrow \Xi$ for Ξ, and a cost function $c : \Xi \rightarrow \Re_{\geqslant 0} \cup \{\infty\}$, such that, for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$:*

$$(2) \quad \gamma(A, \dagger_x) = \begin{cases} \sup(A, u, D(A)) & \textit{if } \sup(A, u, D(A)) \neq \varnothing \\ \sup(A, u, \{p_\Xi\}) & \textit{if } \sup(A, u, D(A)) = \varnothing \textit{ and } c(D(A)) \leqslant u(x) \\ \{\dagger_x\} & \textit{otherwise} \end{cases}$$

Furthermore, Ξ is unique, the restriction of D to $\overline{\Upsilon}$ is unique, and u and c are unique up to the same positive affine transformation.

The confidence ranking, the utility function and the cost function have the standard uniqueness properties. Moreover, the cautiousness coefficient has the sort of uniqueness properties one would expect. Wherever the confidence level appropriate for the choice from a menu counts – wherever there are some costs of deferral for which the decision maker defers and others where he decides – the level of confidence is fixed uniquely by the cautiousness coefficient. Wherever the decision maker chooses no matter the cost of deferral, the precise value of the appropriate level of confidence is not important; all that matters is that it is low enough for the beliefs held to that level of confidence to yield a choice. On these menus, where the precise setting of the appropriate level of confidence does not matter for choice, the cautiousness coefficient is not necessarily unique. In other words, the cautiousness coefficient is unique where the appropriate level of confidence matters, but not where it doesn't.

Remark 3. As noted in Remark 1, the representation can be reformulated in terms of a stakes relation on the set of menus and a cautiousness coefficient that respects this relation, in the style of Hill (2014). This reformulation gives a version of Theorem 1 that yields a stakes relation, as well as the other elements cited. It is straightforward to check that this stakes relation has the same uniqueness properties as the cautiousness coefficient: it is unique on $\overline{\Upsilon}$. To this extent, Theorem 1 provides an independent contribution to the literature on the confidence-based approach set out in Hill (2013, 2014), in the form of a representation result where the stakes relation is endogenised, that is elicited from behaviour rather than assumed. See Section 5 for further discussion.

We note finally that comparative statics analyses can be undertaken on the proposed representation, separating in particular the roles of various elements of the model. Some details are given in Appendix A.

4 Confidence and difference in information as sources of deferral

The representation and characterization of choice in the presence of a costly deferral option lays the ground for a deeper understanding of the relationship between the intrinsic source of deferral explored in this paper and extrinsic sources studied elsewhere in the literature (Section 1.1), as well as for the development of multi-factor theories of deferral. As a preliminary exploration, we focus on a particular, popular model of extrinsic deferral, based on the difference in information between the decision maker and the person who will finally take the decision. We perform a brief comparison with the model proposed in the previous section, and consider how the two aspects could be combined.

4.1 Confidence- versus information-driven deferral

Assume the setup used above, and suppose moreover that if the decision maker – call him DM1 – defers, then the person to whom he defers (who may be his future self) – call him DM2 – will be faced with the same choice, which he cannot defer. A standard simple approach to this sort of decision would be to assume that both decision makers are Bayesian with the same (strictly increasing, zeroed, continuous, affine) utility function $u : \Delta(X) \to \Re$.[17] The information differential between the decision makers can be characterised by a finite set I of signals, with each $i \in I$ associated with an ex post probability measure p_i over S, and a probability measure q over I. DM2 receives one of the signals before deciding, and so has probability measure p_i for the appropriate i. DM1 only has a probability measure over the signal received by DM2, q; his probability measure over states is p where $p(s) = \mathbb{E}_q p_i(s)$ for all $s \in S$.[18] The value of deferral, measured in utility units, is the value of the expected information differential given the cost: that is, the difference between the expected value of choosing ex post after having paid the cost of deferral and the value of choosing ex ante. More precisely, for the choice from a menu A where deferral costs x, it is given by:[19]

$$VI_x(A) = \mathbb{E}_q \max_{h \in A} \mathbb{E}_{p_i} u(h(s) \dot{-} x) - \max_{h \in A} \mathbb{E}_p u(h(s)) \qquad (3)$$

A formulation in monetary units is given by the demand value of deferral (Marschak and Miyasawa, 1968) – the highest monetary amount that DM1 is willing to pay to defer – which, for a menu A, is the (unique) solution z to

[17] The strategic literature on delegation only gets properly started, of course, when the utility functions are different, but we focus on this simple case to highlight the contrast with the model proposed here.

[18] $\mathbb{E}_q$ is the expectation operator given distribution q, defined by $\mathbb{E}_q f(i) = \sum_{i \in I} f(i) dq(i)$.

[19] Recall from Section 2.1 that $h(s) \dot{-} x$ is the result of subtracting the monetary value x from each outcome yielded by the lottery $h(s)$.

$$VI_z(A) = 0 \tag{4}$$

The decision maker (DM1) defers whenever $VI_x(A)$ is positive, or equivalently the cost of deferral is less than the demand value, and decides if not. Hence this approach yields the following representation of the choice correspondence for costly deferral γ: for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$

$$\gamma(A, \dagger_x) = \begin{cases} \sup(A, u, \{p\}) & \text{if } VI_x(A) \leqslant 0 \\ \{\dagger_x\} & \text{otherwise} \end{cases} \tag{5}$$

As concerns the criteria determining whether the decision maker defers or not, there are two evident points of contrast with representation (2). The first concerns the value or cost considerations driving deferral. Under the standard approach, it is the value of the expected information that determines whether deferral is acceptable, whereas under the model proposed in Section 3, the relevant factor is the (psychological) cost of deciding when one would have wanted to defer. Note that, consistently with the distinction between extrinsic and intrinsic reasons for deferring introduced in Section 1.1, the value of information depends on factors beyond DM1's immediate decision – and in particular the information available to the person who will take the decision if it is deferred – whilst the cost of deciding in representation (2) only depends on the decision itself, and in particular the level of confidence appropriate for it. Given the relative nature of the former term, this difference between the two representations does not immediately translate into a perspicuous axiomatic difference in the setup adopted here.

The second point of contrast concerns the existence of a second 'non-cost' criterion determining whether the decision maker defers or not. Under the value of information approach, deferral is entirely dictated by the value considerations just discussed: the decision maker defers whenever the value of the expected information given the cost of deferral is positive. Under representation (2), the decision maker's capacity to decide on the basis of the beliefs in which he has appropriate confidence also plays a role in determining whether he defers or not. In particular, he does not necessarily defer in cases where the psychological cost of deciding outweighs the cost of deferral: whenever he has sufficient confidence in the relevant beliefs, he decides. This difference between the value of information and confidence approaches does have simple behavioral consequences.

Under the value of information approach, if the decision maker defers the decision from a menu A at some non-zero cost of deferral, then for each mixture $A_\alpha h$ of the menu, there is a non-zero cost of deferral at which the decision maker defers the choice. For instance, if the decision maker defers, at some cost, the choice of whether to take a bet where one could win or lose \$1 M, then there is a cost he would pay to defer the choice of whether to bet on the same event but with stakes of \$10.[20] If there is a difference of information between DM1 and DM2 in one case, there will be an information difference in the other. This behavioural pattern may be violated under the model developed in this paper, because, although one might not have enough confidence in the relevant beliefs to decide when \$1 M are at stake, the level of confidence required may be lower when only \$10 are at stake. If one holds the relevant beliefs to that level of confidence, one will decide in the latter decision, no matter the cost of deferral.[21]

We thus conclude that the confidence approach developed in this paper is behaviorally distinguishable from one of the major standard approaches to deferral, or delegation, namely that based on difference or asymmetry of information.[22]

4.2 Combining confidence and information: a simple multi-factor model of deferral

As made clear at the outset, the ultimate aim of this paper is not to claim that there is one 'correct' account of deferral, but rather to draw attention to a different, intrinsic, source of deferral, which eventually may be integrated into existing approaches. To illustrate

[20]The example in the text assumes a linear utility function. As standard, this assumption can be dropped by replacing \$10 by a lottery yielding a 10^{-5} chance of getting \$1 M.

[21]More formally, over and above the independence axiom used here (A3), the information-driven model of deferral (5) satisfies the following condition, whilst the confidence-based model (2) does not:

(Deferral-Independence). For all $A \in \wp(\mathcal{A})$, $h \in \mathcal{A}$, $\alpha \in (0,1]$ and $x \in \Re_{>0}$, if $\dagger_x \in \gamma(A, \dagger_x)$, then there exists $x' \in \Re_{>0}$ such that $\dagger_{x'} \in \gamma(A_\alpha h, \dagger_{x'})$.

[22]A similar analysis applies *mutatis mutandis* to many of the main models in which deferral (leaving a choice open) is based on consideration of the decision maker's future self's possible utility functions (for example, Kreps (1979); Dekel, Lipman, and Rustichini (2001); Gul and Pesendorfer (2001)). Note in particular that the addition of a self-control cost to a representation involving difference in information (including information about future preferences), in the style of Gul and Pesendorfer (2001) for example, would not alleviate the behavioral difference identified above, since the independence axiom in the menu framework that is satisfied by such models essentially implies the condition given in footnote 21. Moreover, in temptation models, the decision maker may decide even if he were paid to defer (ie. if the cost of deferral were negative), whereas this is not a consequence of the confidence-based approach.

its potential interest, we now consider, on a very simple setup, how the two reasons for deferring discussed above – expected difference in information and insufficient confidence – may be combined, and some consequences of the presence of the two.

To this end, consider a decision maker DM3 who perceives the same information structure as DM1 (Section 4.1), but who is also a confidence-based decision maker of the sort proposed in Section 3. More specifically, suppose that, beyond the utility function, the set of signals I, posteriors p_i and the distribution q over I, the decision maker has a centered confidence ranking that is centered on p (DM1's prior probability measure).[23] Whenever DM3 has sufficient confidence given the decision, he chooses, as in representation (2). In these cases, he makes the same choices as DM1. When he lacks sufficient confidence, his behaviour will depend on the value of deferring, which now incorporates confidence- and information-related factors. The value, in utility units, of deferring from menu A when deferral costs x and the decision maker has insufficient confidence in the relevant beliefs becomes:

$$
\text{(6)} \qquad \begin{aligned} VD_x(A) &= \mathbb{E}_q \max_{h \in A} \mathbb{E}_{p_i} u(h(s) \dot{-} x) - \max_{h \in A} \mathbb{E}_p \left(u(h(s)) - c(D(A)) \right) \\ &= VI_x(A) + c(D(A)) \end{aligned}
$$

The first term on the first line is the expected value of the choice made by the person deferred to, which incorporates the cost of deferring x. The second term is the value of deciding rather than deferring for DM3: it is the standard expected value of the best act, but, unlike the equivalent term in (3), it incorporates the psychological cost of deciding from A when one does not hold the relevant beliefs with sufficient confidence for the decision, $c(D(A))$. (Recall that this cost of deciding is in utility units and is normalized to be non-negative.) The second line follows from the first by the definition of VI_x (equation (3)). So DM3 chooses according to a representation that is the same as (2) except that $c(D(A)) \leqslant u(x)$ is replaced by $VI_x(A) + c(D(A)) \leqslant 0$. Axiomatisation of this model and generalizations beyond the terms of this simple analysis are left to future research.

[23] The posteriors p_i could also be centres of centered confidence rankings; however, since ex post the decision maker is forced to choose, he will act like DM2 above, so the rest of the confidence rankings are irrelevant. Likewise, the distribution q over I may be thought of as a centre of a confidence ranking, but we may ignore the rest of the confidence ranking for the purposes of this simple exercise.

The incorporation of confidence into a standard information-based model of deferral thus introduces a second term – the cost of deciding from representation (2) – into the determinant of whether the decision maker defers. Unlike the value of information term, the cost of deciding term does not depend on the information differential between the decision maker and the person who will take the decision if deferred, but only on the decision itself and the level of confidence appropriate for it. It increases the value of deferral, and may thus lead to deferral in situations where the information differential alone would not.

This effect can be clearly seen on the demand value of deferral (that is, the highest amount the decision maker is willing to pay to defer). A decision maker incorporating both information- and confidence-related considerations into his choices, and who is not sufficiently confident in his beliefs given the decision, has a demand value of deferral of y for the choice from a menu A, where:

$$VD_y(A) = VI_y(A) + c(D(A)) = 0 \tag{7}$$

Contrasting this with the demand value under the pure information-based model (given by (4)) reveals that the incorporation of confidence bumps up the amount the decision making is willing to pay to defer. Moreover, this increase is greater for more important decisions, which call for more confidence.[24]

The introduction of confidence into the standard information-based approach may thus be able to explain cases where the information differential between the decision maker and the person to whom he defers does not justify the cost of deferral, but where the decision maker defers nonetheless. For instance, active portfolio management appears to perform too poorly with respect to passive benchmarks to justify its cost (Malkiel, 1995; Gruber, 1996), making it difficult to rationalize delegation solely on the basis of the difference in information (or information-gathering ability) between the investor and the portfolio manager. In such cases, where the decision is important and hence the cost of deciding is high, this cost may drive up the price that the investor is willing to pay to defer, ultimately leading to the choice to delegate.

[24]Basic comparative statics of the demand value of deferral can be immediately read off from equation (7) and the analysis in Appendix A. Decision makers that are more decision averse (see Appendix A for a precise definition of this notion) have higher costs of deciding from a menu A ($c(D(A))$), so their demand value of deferral for A is higher.

Such cases could be just instances of a more general sort of 'calibration error'. The analysis suggests that standard estimations of the information differential based on willingness to pay to delegate or defer tend to be overvalued: part of what may be driving deferral is the reluctance to choose in important decisions on the basis of beliefs in which one has insufficient confidence. Correct estimations of the information differential may need to take this factor into account.

Finally, we note that the introduction of confidence into the information-based approach might also have consequences for what delegation or deferral options the decision maker adopts when he defers. If investors in fact see themselves as buying two services from portfolio managers – information, but also the taking of a difficult decision on their behalf – the relative importance of the two aspects in the decision to delegate may impact upon the competitive pressure on managers to provide high quality information. Exploration of such possible consequences of the introduction of confidence is left for further research.[25]

5 Related literature

Several related papers have already been mentioned previously; we now discuss some points of related literature that have not been treated above.

As noted in the Introduction, Bewley's (1986 / 2002) theory of status quo choice – which can be thought of as deferral to nature – is one of the first axiomatic theories of intrinsic deferral-like phenomena. The essential intuition, which is behind part of the subsequent axiomatic literature on the status quo bias (for example, Masatlioglu and Ok (2005); Ortoleva (2010); Riella and Teper (2014))[26], is that one sticks with the status quo in the absence of appropriate determinate preference. Many discussions of deferral in behavioral

[25] Although the details will of course differ considerably, there is little reason to expect that the general conclusions are specific to information-based approaches. There may even be interesting consequences of integrating confidence into accounts of deferral, or procrastination, based on time inconsistency, which have a different structure and often different primitives. For example, for a single task – such as a single decision to be taken – with a fixed cost, procrastination tends to decrease with increases in the per-period benefit (O'Donoghue and Rabin, 2001), suggesting that there would be less deferral in important decisions. Adding confidence to the model would imply that the cost (which is exogenous in the cited model) increases with the importance of the decision, tempering the mitigating effect of increased importance on deferral. This may contribute to explaining procrastination in important decisions such as the choice of retirement policy.

[26] A technically related paper drawing on a different motivation is Savochkin (2014).

economics (for example, Tversky and Shafir (1992)) and marketing (for example, Dhar (1997)), though generally couched in the context of decision under certainty, make reference at times to apparently similar intuitions that deferral results from preferences being incompletely constructed or 'conflicting'. To the extent that it is the decision maker's actual preferences that drive deferral, these are involved with intrinsic deferral. Although the theory proposed in this paper may be interpreted as a theory of status quo choice, the interpretation is more general, covering and focussing on cases of deferral to one's later self, and deferral to someone else.

As suggested previously, there is a relationship between representation (2) and incomplete preferences. Rather than discussing the whole literature on incomplete preferences, we focus on several conceptually and formally related papers.[27] Danan (2003b) and Kopylov (2009) both involve incomplete preference relations where incompleteness is interpreted in terms of postponing a decision. The former paper cashes out this option in terms of preference for flexibility and hence, as discussed in Sections 1.1 and 4.1, is concerned with extrinsic deferral (see Danan (2003a) for a related study in a choice-theoretical framework). The latter paper, which uses the unanimity representation à la Bewley and is noncommittal about the relationship between postponing a decision and preference for flexibility, may be interpreted in terms of intrinsic deferral. However, it does not offer an account of behaviour when deferral or postponement is costly.

The closest paper in this literature is Hill (2014), which, though it does not involve a deferral interpretation of incompleteness, can technically be thought of as a special case of (2) where deferral is free and all menus are binary. However, like Hill (2013), it assumes an exogenously-given notion of the stakes involved in, or importance of, a decision. This assumption is objectionable, especially if one considers the importance of a decision to be a subjective judgement on the part of the decision maker. No such assumption is made here. Indeed, as discussed in Remarks 1 and 3, a notion of stakes can be defined from the endogenous elements of representation (2). Theorem 1 can thus be thought of as providing a behavioural foundation for the notion of stakes assumed in the cited papers, and hence an answer to the aforementioned criticism of the confidence-based approach that they set out. Indeed, this could be considered a separate contribution of the theorem; it is, to our knowledge, the first behavioural foundation for the notion of stakes in the literature.

[27] Readers seeking a more developed discussion of this literature are referred to Hill (2014).

The cost of deciding involved in representation (2) is reminiscent of cost factors involved in representations studied by Ergin and Sarver (2010); Ortoleva (2013); Buturak and Evren (2014). In a menu-preference framework, Ergin and Sarver (2010) derive endogenously the elements of a costly information acquisition representation, where different information structures have costs, called ‘contemplation costs’. Given the conceptual proximity to the literature on information acquisition and preference for flexibility, the general remarks concerning the relationship between the approach taken here and standard approaches (see Sections 1.1 and 4.1) continue to apply. In particular, when applied to the question of deferral, Ergin and Sarver (2010) evidently take an extrinsic approach, which, even ignoring technical differences and differences in framework, clearly yields a different representation from that proposed here.

Buturak and Evren (2014), working in a choice-theoretic framework with a fixed ‘default option’ that is always present and no explicit cost of taking that option, introduce several related models incorporating a value of this option. The default option in their models can be either interpreted as a (fixed) status quo or as the opportunity to postpone the decision. In the most structured ‘subjective state space’ representation, both interpretations involve consideration of the decision maker’s ex ante uncertainty about his future tastes. Hence, when interpreted in terms of deferral, this representation adopts an extrinsic approach, as opposed to the intrinsic approach taken here. The less structured ‘variable threshold’ representation, under the status quo interpretation of the default option, is related to the status quo literature mentioned above, and many of the points made there apply to this case. Behavioural differences between these representations and our representation (2) include the fact that the former demand that the decision maker defers from some singleton menus, whereas the latter implies that the decision maker never defers from singleton menus. Moreover, once this aspect is set aside, the independence axiom involved in the former representations appears to be stronger than our A3, to the extent that it demands that the decision maker choose from a mixture of a menu whenever he chooses from the menu itself (see also the discussion in Section 3.2.1).[28]

[28]More precisely, in their models, the decision maker chooses from the mixture of a menu with an alternative if he chooses from the menu and from the singleton set containing the alternative; as noticed, this latter condition is automatically satisfied in the model proposed here. Another way of seeing the difference is by noting that, though our representation (2) can be written in the same form as their VT representation (their (1)), the threshold function will not be affine, as assumed in their representation (Definition 1).

Ortoleva (2013) proposes a representation of preferences over menus where the valuation of each menu differs from the standard 'additive EU representation' (Kreps, 1979; Dekel, Lipman, and Rustichini, 2001) by the addition of a 'thinking cost' that is a function of the menu. Despite the difference in interpretation – the thinking cost is generally motivated in terms of computational constraints, rather than the importance of the decision – the cost of deciding involved in representation (2) could be thought of as a possible source of Ortoleva's thinking cost, insofar as it represents the psychological cost of taking certain decisions. However, even putting aside technical differences and differences in framework, the representation in Ortoleva (2013) is quite different from ours, because it is concerned with a different choice problem. Ortoleva considers the problem of deciding which choice the decision maker would like to face at some future moment – and the cost of deciding or thinking in the choice he eventually faces will be a factor in this ex ante decision, though it does not affect the ex post choice taken. By contrast, we consider behavior in a single choice situation where deferral is an option – and the cost of deciding may influence the decision to defer or not.[29]

6 Conclusion

One potential source of deferral is lack of confidence in the beliefs needed to decide. Whilst recognised in some studies of deferral behavior – such as deferral to nature or to the status quo – this source is completely absent from many accounts of deferral – in particular standard accounts of delegation to others or deferral to one's future self. This paper makes a start at incorporating it into our models of such cases of deferral.

We develop a theory of deferral in which the decision maker's confidence in his beliefs plays a double role. It drives deferral, insofar as the decision maker is inclined to defer only

[29]This difference can be seen in the representations themselves, by considering the choice between f and g in the presence of a free deferral option: written in menu language, this is the choice among $\{f\}$, $\{g\}$ and $\{f, g\}$. Under Ortoleva's representation, there is no cost of thinking involved in the valuation of $\{f\}$ and $\{g\}$, but there may be a positive cost attached to $\{f, g\}$; under our representation, opting for $\{f\}$ or $\{g\}$ incurs a cost of deciding, whereas deferral – $\{f, g\}$ – incurs no cost. This is due to the difference in the choice problems considered. Since Ortoleva is considering what choice the decision maker would like to be faced with, if he is faced with $\{f\}$, no thinking will be required, so the cost is zero. Since we are considering behaviour in the choice between f and g with the possibility of deferring, $\{f\}$ amounts to making a decision and hence may incur a cost.

when he lacks sufficient confidence in the relevant beliefs given the decision to be taken; whenever he has sufficient confidence, he decides. It prices deferral, to the extent that the value of the deferral option is determined by the confidence level appropriate for the decision under consideration. The theory applies naturally to situations where deferral is costly, which are the norm in cases of deferral to others or to one's future self. In these situations, a decision maker represented by the model decides if he has sufficient confidence in the relevant beliefs given the decision; if he lacks sufficient confidence, he decides whenever the cost of deferral outweighs its value, as determined by the confidence level appropriate for the decision, and he defers if not.

A behavioural axiomatization of this model is given, in which all the elements are derived endogenously. Choice-theoretically, the theory is quite reasonable: the essential axiomatic difference from the standard Savagean theory of decision under uncertainty (applied in a choice-theoretic setting) is to weaken the equivalent of the preference-theoretic completeness, transitivity and independence axioms to allow the decision maker to defer in cases where the standard axioms would have demanded decision.

The formulation of the model and its axiomatic analysis clarifies the relationship between confidence as a source of deferral and standard mechanisms for deferral to others or to one's future self considered in the literature, such as those based on expected difference or asymmetry of information between the deferrer and the deferree. In particular, it reveals behavioural patterns exhibited by information-based models but not by the proposed confidence-based one, hence showing that the two approaches to deferral are behaviourally distinct.

Finally, the proposed model provides a starting point for the incorporation of confidence into standard approaches to deferral. A straightforward integration of confidence into an information-differential model yields a simple additive formula for the effect of confidence on the value of deferral. The model suggests that the introduction of confidence as a source of deferral increases the price the decision maker is willing to pay to defer, and that this effect is stronger for more important decisions (which call for more confidence). This may help explain some cases that are problematic for the standard approach: although the expected information differential may be insufficient to explain the observed willingness to pay for deferral or delegation, lack of confidence in the relevant beliefs, combined with the importance of the decision, would imply deferral even at relatively high costs.

Appendix A Attitudes to decision and deferral

In this Appendix, we perform some basic comparative statics on representation (2), bringing out the role of the various elements in the decision maker's attitude to deciding or deferring.

In the current context, there is a particularly simple comparison of decision makers' attitude to deciding: if decision maker 2 decides, rather than deferring, in every choice situation in which decision maker 1 decides, then decision maker 2 has less trouble deciding, and seeks deferral less, than decision maker 1. We shall say that decision maker 2 is less *decision averse*, or equivalently less *deferral seeking*, than decision maker 1. Formally, we say that (a decision maker whose choice correspondence is) γ^1 is *more decision averse* than γ^2 when, for all $A \in \wp(\mathcal{A})$, $x \in \Re_{\geqslant 0}$, if $\dagger_x \notin \gamma^1(A, \dagger_x)$ then $\dagger_x \notin \gamma^2(A, \dagger_x)$.

Beyond this benchmark comparison, we consider two others, which shall be useful in the separation of the roles of the different elements of the model. The first involves the comparison of when the cost of deferral is motivating. Recall from the discussion in Section 3.2 that a cost is said to be motivating for a menu when it can drive the decision from it: there is a version of the menu from which the decision maker decides when deferral has this cost, though he does not decide when deferral is free. We shall say that decision maker 2 is more *cost motivated* than decision maker 1 if whenever a cost is motivating for a menu for decision maker 1, it is motivating for decision maker 2. Formally, we say that (a decision maker whose choice correspondence is) γ^1 is *less cost motivated* than γ^2 when, for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$, if $\overline{\gamma^1}^x(A) = \varnothing$ then $\overline{\gamma^2}^x(A) = \varnothing$.[30]

Decision aversion compares, for a given menu, the costs of deferral at which the decision makers decide. Another possible comparison focusses on the menus for which the cost is motivating, rather than the cost itself. If a decision maker decides from a menu A at every cost which is motivating for some other menu B, then this provides an indication about his deferral behaviour, relative to the extent to which cost motivates his choices. The more menus B for which this holds, the more he is inclined to decide rather than defer from A. If whenever decision maker 1 chooses from a menu at every cost that is motivating for another menu, decision maker 2 does the same, then there is a sense in which decision maker 2 exhibits less decision aversion. We shall say that decision maker 2 is less

[30] Recall from the discussion in Section 3.2 that, for a choice correspondence for costly deferral γ, $\overline{\gamma}^x(A) = \varnothing$ exactly when x is motivating for A.

motivation-calibrated decision averse than decision maker 1; the terminology reflects the fact that decision maker's cost motivation (rather than the cost itself) is used as a yardstick in the comparison. Formally, we say that (a decision maker whose choice correspondence is) γ^1 is *more motivation-calibrated decision averse* than γ^2 when, for all $A, B \in \mathcal{A}$, if $\dagger_x \notin \gamma_1(A, \dagger_x)$ for all $x \in \Re_{\geqslant 0}$ such that $\overline{\gamma^1}^x(B) = \varnothing$, then $\dagger_x \notin \gamma_2(A, \dagger_x)$ for all $x \in \Re_{\geqslant 0}$ such that $\overline{\gamma^2}^x(B) = \varnothing$.

The weakness of these notions – none of them ask that the decision makers choose the same acts when they decide – mean that the implications for the relationships between the elements of the representation are also weak. For a finer analysis, we shall follow a standard strategy in the literature and make assumptions sufficiently strong to guarantee that the decision makers have the same utilities and beliefs. To this end, we introduce the following notion.

Definition 3. Let γ satisfy axioms A1–A10. The *confidence-in-choice* relation $\preccurlyeq$ on $\mathcal{A} \times \wp(\mathcal{A})$ is defined as follows: for any $A, B \in \wp(\mathcal{A})$ and $f, g \in \mathcal{A}$, $(f, A) \preccurlyeq (g, B)$ iff, for all $x \in \Re_{\geqslant 0}$, if $f \in \overline{\gamma}^x(A)$, then $g \in \overline{\gamma}^x(B)$.

Recall from Section 3.2 that a cost is not motivating for a menu if, whenever the decision maker decides at this cost, he would decide even if deferral was free: he is, so to speak, sure enough to choose even if there were no cost to deferral. So if a cost x is not motivating for the choice from B and g is chosen, but x is motivating for the choice from A, then this can be taken as an indication that the decision maker is *more confident in his choice of g from B than in his choice from A*. Definition 3 introduces a confidence-in-choice relation that, in these situations, ranks the choice of g from B ((g, B)) higher than any choice from A (for example, (f, A)).

We shall say that two decision makers are confidence equivalent if they have the same confidence-in-choice relation.

Definition 4. Let γ^1 and γ^2 satisfy axioms A1–A10. γ^1 and γ^2 are *confidence equivalent* if $\preccurlyeq^1 = \preccurlyeq^2$.

Proposition A.1. *Let γ^1 and γ^2 satisfy axioms A1–A10, and be represented by (u_1, Ξ_1, D_1, c_1) and (u_2, Ξ_2, D_2, c_2) respectively. γ^1 and γ^2 are confidence equivalent if and only if u_2 is a positive affine transformation of u_1 and $\Xi_1 = \Xi_2$.*

This proposition confirms that a decision maker's confidence in his choices is entirely determined by his utilities and his confidence in beliefs. This is to be expected: to the extent that his choices are dictated by his utilities and his beliefs, it is reasonable that confidence in choices 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. Once differences in utilities and confidence in beliefs are accounted for, by comparing decision makers who have the same confidence in choices, the different comparisons mentioned above are entirely characterised by the relationship between the cautiousness coefficients and the cost functions, as the following proposition shows.

Proposition A.2. *Let γ^1 and γ^2 satisfy axioms A1–A10 and be confidence equivalent. Then:*

(i) *γ^1 is more decision averse than γ^2 if and only if there exist representations (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) of γ^1 and γ^2 respectively such that $c_2(D_2(A)) \leqslant c_1(D_1(A))$ for all $A \in \overline{\Upsilon_2}$ and $D_2(A) \subseteq D_1(A)$ for all $A \notin \overline{\Upsilon}_2$.*

(ii) *γ^1 is less cost motivated than γ^2 if and only if there exist representations (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) of γ^1 and γ^2 respectively such that $c_2(\mathcal{C}) \leqslant c_1(\mathcal{C})$ for all $\mathcal{C} \in \Xi$.*

(iii) *γ^1 is more motivation-calibrated decision averse than γ^2 if and only if there exist representations (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) of γ^1 and γ^2 respectively such that $D_2(A) \subseteq D_1(A)$ for all $A \in \wp(\mathcal{A})$.*

Since the cost function is unique, given the utility function, the comparison in part (ii) holds for all representations using the same utility function. Similarly, since the cautiousness coefficient D is only unique on $\overline{\Upsilon}$, it really only makes sense to compare D_1 and D_2 on $\overline{\Upsilon_1} \cap \overline{\Upsilon_2}$, and the comparisons on these sets (in parts (i) and (iii)) hold for all representations using the same utility function. (Note that both notions of decision aversion in fact imply that $\overline{\Upsilon_1} \supseteq \overline{\Upsilon_2}$, as is clear from the proof.) The formulation used for part (iii) in particular emphasises that the restriction to $\overline{\Upsilon_2}$ in the uniqueness of the cautiousness coefficient does not complicate applications: insofar as the behavioural consequences of comparisons of motivation-calibrated decision aversion are concerned, one can assume that the cautiousness coefficient containment condition holds everywhere.

Proposition A.2 brings out both the characterisation of the standard notion of decision aversion and the behavioural effects of the cautiousness coefficient and the cost function. More decision averse decision makers assign a higher cost to (the confidence level associated with) every menu. For each menu, they value the option of deferring from that menu more than less decision averse counterparts. Whilst not surprising, this does not separate the role of the cautiousness coefficient and the cost function, insofar as the cost assigned to a menu depends both on the level of confidence deemed appropriate for choice from the menu (determined by the cautiousness coefficient) and the cost of deciding when this level of confidence is required (determined by the cost function).

The other two parts of the result provide this separation. Comparison in terms of the cost function alone corresponds precisely to a difference in cost motivation: decision makers who are less cost motivated have a higher cost of deciding for every confidence level. It is the degree to which certain costs motivate decision makers' decisions that reveals the relationship between their costs of deciding at given confidence levels. By contrast, ordering in terms of motivation-calibrated decision aversion is characterised precisely in terms of an appropriate ordering of the cautiousness coefficients. It is intuitive that decision makers who assign higher levels of confidence to each menu are more decision averse: part (iii) of Proposition A.2 shows that this is the case, when decision aversion is measured on the motivation rather than the cost yardstick. Note that these results confirm the interpretations adopted for the elements of the model. Confidence drives deferral – and the cautiousness coefficient reflects tastes for choosing on the basis of limited confidence, when measured on a scale that is cost independent (the scale provided by motivation). Confidence also prices deferral – and the cost function captures the value of deferring, or equivalently the cost of deciding, fleshed out entirely in terms of the sensitivity of deferral behaviour to the cost of deferral.

Naturally, piecing together parts (ii) and (iii), one obtains a behavioural characterisation of when two decision makers satisfy both the containment condition on cautiousness coefficients and the ordering condition on cost functions: one is both less cost motivated and more motivation-calibrated decision averse than the other. It is evident from the proposition that these two comparisons imply that the decision maker is more decision averse. As Proposition C.3 in Appendix C shows, apart from this implication, the comparisons are independent: no other pair of comparisons imply the third.

Appendix B Proof of Theorem 1

Throughout the remaining Appendices, B will denote the space of real-valued functions on S, and $ba(S)$ will denote the set of additive real-valued set functions on S, both under the Euclidean topology. B is equipped with the standard order: $a \leqslant b$ iff $a(s) \leqslant b(s)$ for all $s \in S$. For $x \in \Re$, let x^* be the constant function taking value x.

The main part of Theorem 1 consists in showing the sufficiency of the axioms ((i) to (ii) direction), the proof of which proceeds as follows. We first construct a binary relation $\leqq$ on $\wp(\mathcal{A})$ satisfying appropriate properties. This essentially orders menus according to whether a higher or lower confidence level is appropriate. The essential idea of the construction is that, for $A, B \in \Upsilon$, $A \leqq B$ iff $\iota(A) \leqslant \iota(B)$; the main work in the construction (Lemma B.1) is to extend this definition to the whole of $\wp(\mathcal{A})$ in such a way as the retain the appropriate properties. Then we establish (Lemma B.2) the representation for the case where deferral is free. To this end, for each indifference class r under $\leqq$ – which, recall, can be thought of as corresponding to a level of confidence – we define a function γ_r from menus to (perhaps empty) subsets, which can be thought of as representing the choices from menus considered 'as if' the appropriate level of confidence was r. We show (Lemmas B.4 to B.6) that for every non-minimal r there is a closed convex set of probability measures $\mathcal{C}_r$ representing γ_r in the sense that the elements selected by γ_r are the optimal ones according to the unanimity rule applied with $\mathcal{C}_r$. Lemmas B.8 and B.9 show that the choice correspondence for minimal r can be represented according to the unanimity rule with the intersection of the $\mathcal{C}_r$ for the other stakes levels, and that this set is a singleton. By Lemma B.7, the $\mathcal{C}_r$ form a nested family of sets, and we thus have a confidence ranking. By Lemmas C.8 and C.9, this confidence ranking is continuous, and by Lemma C.10, it is strict. By construction, the function that assigns to any level r the set $\mathcal{C}_r$ is a well-defined cautiousness coefficient. Moreover, the function that assigns to any set $\mathcal{C}_r$ in the confidence ranking the utility of $\iota(A)$ for any $A \in r$ is a well-defined cost function. Finally we show that these elements correctly capture the part of the representation where deferral has non-zero cost (Lemma B.10). We detail the main steps below; proofs of the technical lemmas are relegated to Appendix C.

B.1 Sufficiency of axioms

Note firstly that since $X = \Re$ is a separable metric space, the space of all Borel probability measures on X, and hence $\Delta(X)$ (which, recall from Section 2.1, is the subspace of finitely additive probability measures on X), can be equipped with a separable metric (Billingsley, 2009, p72; Aliprantis and Border, 2007, Theorem 15.12). Hence, $\mathcal{A}$ is a separable metric space, equipped with the product metric, and $\wp(\mathcal{A})$ is a metric space, under the Hausdorff metric. If $\Phi = \varnothing$, then by A4, $\dagger_x \notin \gamma(A, \dagger_x)$ for all $A \in \wp(\mathcal{A})$, $x \in \Re_{\geqslant 0}$, and the result follows from the standard representation theorem for expected utility. Throughout the rest of the proof, we thus assume that $\Phi \neq \varnothing$. We begin by stating some preliminary results (see Appendix C for proofs).

The first is a generalisation of Hill (2012, Theorem 2), which held for finite menus, to the case of infinite menus (with topological structure).

Theorem 2. *Let X be a metric space and $\wp(X)$ the set of non-empty compact subsets of X, endowed with the Hausdorff topology. Let $\gamma : \wp(X) \to 2^X$ be a function such that $\gamma(A) \subseteq A$ for all $A \in \wp(X)$. Suppose that γ satisfies the following continuity property: for all sequences $A_i \in \wp(X)$ and $x_i \in X$ with $A_i \to A$ and $x_i \to x$, if $x_i \in \gamma(A_i)$ for all i, then $x \in \gamma(A)$.*[31] *The following are equivalent:*

(i) *There exists a reflexive, transitive binary relation $\preceq$ such that $\gamma(A) = \{x \in A \mid x \succeq y \ \forall y \in A\}$;*

(ii) *γ satisfies the following properties:*

α	*if $x \in A \subseteq B$ and $x \in \gamma(B)$, then $x \in \gamma(A)$*
π	*if $x \in A$, $y \in A \cap B$, $y \in \gamma(B)$ and $x \in \gamma(A)$, then $x \in \gamma(A \cup B)$*
$sing$	*if $A = \{x\}$, then $x \in \gamma(A)$.*

Moreover, $\preceq$ is unique.

Finally, if γ always takes non-empty values, (ii) is equivalent to the existence of a reflexive, transitive and complete preference relation representing γ, even in the absence of the continuity assumption.

[31] Note that γ is not necessarily a correspondence, because it is not necessarily non-empty-valued. The continuity condition, which is essentially a version of the standard upper hemi-continuity condition for correspondences, is thus stated explicitly.

Secondly, Lemma C.1 establishes that there is a strictly increasing zeroed continuous affine utility function $u : \Delta(X) \to \Re$ representing the restriction of γ to menus of constant acts. We fix such a u for use throughout the proof, and let $K = u(\Delta(X))$ and $B(K)$ be the set of functions in B taking values in K. Note that since u is strictly increasing, K is an open interval. $B(K)$ is naturally isomorphic to a subset of $\Re^{|S|}$, and we take it to be equipped with the Euclidean topology. Lemma C.2 guarantees that, under u, γ generates a well-defined choice correspondence for costly deferral on $\wp(B(K)) \times \mathcal{D}$, which, with slight abuse of notation, we denote by γ.

B.1.1 Construction of 'higher level of confidence appropriate' relation

Lemma B.1. *There exists a non-trivial, continuous weak order*[32] $\leqq$ *on* $\wp(\mathcal{A})$ *satisfying:*

i. *for all* $A, B \in \Upsilon$, $A \leqq B$ *iff* $\iota(A) \leqslant \iota(B)$*;*

ii. *for all* $A \in \Upsilon$, $B \in \Phi \backslash \Upsilon$, *if* $\iota(A) > mmc(B)$, *then* $A > B$*;*

iii. *for all* $A, B, \in \wp(\mathcal{A})$, *if* $A \overset{e.e.}{\simeq} B$, *then* $A \equiv B$*;*

iv. *for all* $A, B \in \Phi \subseteq \wp(\mathcal{A})$, *there exists* $h, h' \in \mathcal{A}$ *and* $\alpha, \alpha' \in (0, 1]$ *such that* $A_\alpha h \leqq B \leqq A_{\alpha'} h'$*;*

v. *there exists* $A \in \Phi$ *such that, for all* $B \in \wp(\mathcal{A})$, $A \leqq B$.

where $\equiv$ *and* $<$ *are the symmetric and asymmetric parts of* $\leqq$ *respectively.*

Proof. Recall that, by assumption, Φ is non-empty. Note that, by A7, $\wp(\Delta(X)) \subseteq \Phi^c$. We first define the function $\tilde{\iota} : \Phi \to \Re$ as follows: $\tilde{\iota}(A) = \iota(A)$ if $A \in \Upsilon$ and $\tilde{\iota}(A) = mmc(A)$ otherwise. By A9 part iii. $\tilde{\iota}$ is continuous.

By Lemma C.1, u is a strictly increasing zeroed continuous affine utility function representing the restriction of γ to sets of constant acts. Note that $\{d \in \Delta(X) \mid \gamma(\{c, d\}, \dagger_0) = \{c\}\}$ and $\{d \in \Delta(X) \mid \gamma(\{c, d\}, \dagger_0) = \{d\}\}$ is a subbase for the set of open subsets of $\Delta(X)$ that are pre-images of sets in K under u; since the image under u of these sets is $\{y \in K \mid y < u(c)\}$ and $\{y \in K \mid y > u(c)\}$, a subbase for the set of open subsets of K, u is a quotient map (Hart, Nagata, and Vaughan, 2004, Section b-4). Moreover, there

[32] A weak order is a reflexive, transitive and complete binary relation; a relation $\leqslant$ on a topological space X is continuous if $\{y \in X \mid x \leqslant y\}$ and $\{y \in X \mid x \geqslant y\}$ are closed in X.

is a many-to-one map $\hat{u} : \mathcal{A} \to B(K)$, given by $\hat{u}(f)(s) = u \circ f(s)$, for all $f \in \mathcal{A}$, $s \in S$. Since u is a quotient map and $\hat{u}$ is a finite product of quotient maps, $\hat{u}$ is also a quotient map. Let $\tilde{u} : \wp(\mathcal{A}) \to \wp(B(K))$ be the generated map between the compact subsets of $\mathcal{A}$ and $B(K)$ (since $\hat{u}$ is continuous, the image of each element of $\wp(\mathcal{A})$ is compact). By Aliprantis and Border (2007, Theorem 3.91), the Hausdorff topology on $\wp(\mathcal{A})$ coincides with the Vietoris topology on $\wp(\mathcal{A})$, and similarly for $\wp(B(K))$. The collection of sets of the form $G_0^h \cap G_1^l \cap \cdots \cap G_n^l$, where the G_i are open subsets of $B(K)$, $G^h = \{F \in \wp(B(K)) \mid F \subset G\}$ and $G^l = \{F \in \wp(B(K)) \mid F \cap G \neq \varnothing\}$, is a base for the Vietoris topology on $\wp(B(K))$ (Aliprantis and Border, 2007, Definition 3.89). So $H \subseteq \wp(B(K))$ is open iff it is the union of such sets, and this holds iff $\tilde{u}^{-1}(H)$ is the union of sets of the form $\tilde{u}^{-1}(G_0^h) \cap \tilde{u}^{-1}(G_1^l) \cap \cdots \cap \tilde{u}^{-1}(G_n^l)$ with G_i open in $B(K)$. But since $\tilde{u}^{-1}(G^h) = \{\tilde{u}^{-1}(F) \mid F \in \wp(B(K)),\ F \subset G\} = \bigcup_{G' \in \hat{u}^{-1}(G)} \{F \in \wp(\mathcal{A}) \mid F \subset G'\} = \bigcup_{G' \in \hat{u}^{-1}(G)} (G')^h$ and $\tilde{u}^{-1}(G^l) = \{\tilde{u}^{-1}(F) \mid F \in \wp(B(K)),\ F \cap G \neq \varnothing\} = \bigcup_{G' \in \hat{u}^{-1}(G)} \{F \in \wp(\mathcal{A}) \mid F \cap G' \neq \varnothing\} = \bigcup_{G' \in \hat{u}^{-1}(G)} (G')^l$, $\tilde{u}^{-1}(H)$ is the union of sets of the form $\tilde{u}^{-1}(G_0^h) \cap \tilde{u}^{-1}(G_1^l) \cap \cdots \cap \tilde{u}^{-1}(G_n^l)$ with G_i open in $B(K)$ iff it is a pre-image of an open subset of $\wp(B(K))$ under $\tilde{u}$ that is the union of sets of the form $(G'_0)^h \cap (G'_1)^l \cap \cdots \cap (G'_n)^l$ where the G'_i are open subsets of $\mathcal{A}$. By the definition of the Vietoris topology on $\wp(\mathcal{A})$, this latter property holds iff $\tilde{u}^{-1}(H)$ is open. So H is open in $\wp(B(K))$ iff $\tilde{u}^{-1}(H)$ is open in $\wp(\mathcal{A})$. Hence $\tilde{u}$ is a quotient map between $\wp(\mathcal{A})$ and $\wp(B(K))$ taken with the Vietoris, or equivalently Hausdorff, topology.

Recall that, with slight abuse of notation, we use $\gamma : \wp(B(K)) \times \mathcal{D} \to 2^{B(K) \cup \mathcal{D}} \backslash \varnothing$ to denote the choice correspondence for costly deferral generated on $\wp(B(K)) \times \mathcal{D}$ by γ under $\tilde{u}$; similarly, we use ι to denote the image of ι and likewise for mmc, Υ, Φ and $\tilde{\iota}$. (Lemma C.2 and A6 imply that γ is well-defined on $B(K)$ and that ι, mmc, $\tilde{\iota}$, Υ and Φ are well-defined on $\wp(B(K))$.) Since $\tilde{u}$ is a quotient map, it follows from a standard result (Kelley, 1975, Theorem 3.9) that the continuity of $\tilde{\iota}$ implies that the function it generates on $\Phi \subseteq \wp(B(K))$ is continuous.

Let $\underline{x} = \inf \iota(\Upsilon)$; by the definition of $\tilde{\iota}$, $\inf \tilde{\iota}(\Upsilon) = \inf \tilde{\iota}(\Phi) = \underline{x}$. For each $A \in \Phi$ such that $mmc(A) \neq \underline{x}$, $\{\alpha A + (1-\alpha) l \mid \alpha \in (0,1], l \in B(K)\} \cap (\overline{\Upsilon})^c \neq \varnothing$, by A10 part ii. For each such A, let $\widehat{A}$ be any element of $\{\alpha A + (1-\alpha) l \mid \alpha \in (0,1], l \in B(K)\}$ such that $d(\hat{A}, \overline{\Upsilon}) > \frac{1}{2} \sup_{A' \in \{\alpha A + (1-\alpha) l \mid \alpha \in (0,1], l \in B(K)\}} d(A', \overline{\Upsilon})$ (where d is the Euclidean metric on $B(K)$). Let $\mathcal{B} = \overline{\{\widehat{A} \mid A \in \wp(B(K)),\ mmc(A) \neq \underline{x}\}}$. Note that, by A9 part i., $A \notin \Upsilon$ for all $A \in \mathcal{B}$. We now show that, if $\mathcal{B} \cap \overline{\Upsilon} \neq \emptyset$, then $\tilde{\iota}(\mathcal{B} \cap \overline{\Upsilon}) = \underline{x}$. For any $A \in \mathcal{B} \cap \overline{\Upsilon}$, the definition of $\mathcal{B}$ and the continuity of the metric implies that there is no $\alpha A + (1-\alpha)l$, for $\alpha \in (0,1]$, $l \in B(K)$, with distance from $\overline{\Upsilon}$ strictly greater than the distance of A from $\overline{\Upsilon}$, whence, by A10 part ii., $mmc(A) = \underline{x}$. Since $A \notin \Upsilon$, it follows from the definition of $\tilde{\iota}$ that $\tilde{\iota}(A) = mmc(A) = \underline{x}$, as required.

Now consider the correspondence $t : \Phi \rightrightarrows \Re$ (ie. map from $\Phi \to 2^{\Re}\backslash\emptyset$) defined as follows: $t(A) = \{\tilde{\iota}(A)\}$ if $A \in \overline{\Upsilon}$; $t(A) = \{\underline{x}\}$ if $A \in \mathcal{B}$; and $t(A) = [\underline{x}, \tilde{\iota}(A)]$ otherwise. t is obviously a well-defined closed and convex-valued correspondence. We now show that it is lower hemicontinuous. By Aliprantis and Border (2007, Theorem 17.21), it suffices to show that for any $(A_n) \in \Phi$ with $A_n \to A$, and any $y \in t(A)$, there exists $y_n \to y$ with $y_n \in t(A_n)$ for all n. Whenever $A \notin \mathcal{B} \cap \overline{\Upsilon}$, this is an immediate consequence of the continuity of $\tilde{\iota}$, the fact that $\overline{\Upsilon}$ and $\mathcal{B}$ are closed, and of the definition of t. Whenever $A \in \mathcal{B} \cap \overline{\Upsilon}$, as shown above, $\tilde{\iota}(A) = \underline{x}$, so the continuity of t at A follows from the continuity of $\tilde{\iota}$. Hence t is lower hemicontinuous.

Since Φ is metric, and hence paracompact, the Michael Selection Theorem (Aliprantis and Border, 2007, Theorem 17.66) implies that there exists a continuous function $s : \Phi \to \Re$ such that $s(A) \in t(A)$ for all $A \in \Phi$. By definition s agrees with $\tilde{\iota}$ on $\overline{\Upsilon}$, and it attains its infimal value.

It remains to extend s to $B(K)$. Since $\wp(B(K))$ is a metric space, so is Φ. Since s is continuous on this space, by Aliprantis and Border (2007, Lemma 3.12) there exists an equivalent metric on Φ under which s is Lipschitz, and hence uniformly continuous. By Aliprantis and Border (2007, Lemma 3.11), there is a uniformly continuous extension of s to the closure of Φ under this metric, which, since the metrics are equivalent, coincides with $\overline{\Phi}$, the closure under the Hausdorff metric. Let $\hat{s}$ be this extension; it is a continuous function (under the Hausdorff metric on $\wp(B(K))$) on $\overline{\Phi}$. Moreover, by definition, $\hat{s}(\overline{\Phi}) = \overline{s(\Phi)}$.

Since $\wp(B(K))$ is a metric space (and hence normal) and $\overline{\Phi}$ is closed, Tietze's extension theorem (Aliprantis and Border, 2007, Theorem 2.47) applies, and implies that there is a continuous real-valued function σ on $\wp(B(K))$ and agreeing with $\hat{s}$ on $\overline{\Phi}$. Moreover, σ can be chosen such that $\sigma(\wp(B(K))) \subseteq \overline{\hat{s}(\overline{\Phi})} = \overline{s(\Phi)}$.

Define the relation $\leqq$ on $\wp(\mathcal{A})$ as follows: for any $A, B \in \wp(\mathcal{A})$, $A \leqq B$ iff $\sigma(\tilde{u}(A)) \leqslant \sigma(\tilde{u}(B))$. Since σ is a continuous real-valued function defined everywhere on $\wp(B(K))$,

$\leqq$ is a continuous weak order; since σ is defined on $\wp(B(K))$, $\leqq$ satisfies property iii. By definition, s and hence σ is not a constant function, so $\leqq$ is non-trivial. Moreover, s attains its infimal value, so σ attains its infimal value with an element in Φ, and hence $\leqq$ satsifies property v. We now show that it satisfies the other required properties.

Property i. follows immediately from the fact that, for all $A \in \wp(\mathcal{A})$ with $A \in \Upsilon$, by definition, $\sigma(\tilde{u}(A)) = \iota(A)$. As concerns property ii., if $A \in \Upsilon$ and $B \in \Phi \backslash \Upsilon$, then, by the definition of σ, $\sigma(\tilde{u}(A)) = \iota(A)$ and $\sigma(\tilde{u}(B)) \leqslant mmc(B)$. Hence if $\iota(A) > mmc(B)$, then $\sigma(\tilde{u}(A)) > \sigma(\tilde{u}(B))$, as required. We now show that $\leqq$ satisfies property iv. On the one hand, by A10 part i., for any $A, B \in \Phi$ with $\sigma(\tilde{u}(A)) < \sigma(\tilde{u}(B)) = x$, there exists $\alpha' \in (0,1]$ and $h' \in \mathcal{A}$ such that $\sigma(\tilde{u}(A_{\alpha'}h')) = \iota(A_{\alpha'}h') \geqslant x$, as required. On the other hand, to treat the case where $A, B \in \Phi$ with $\sigma(\tilde{u}(A)) > \sigma(\tilde{u}(B)) = x$, we consider two cases. If $mmc(A) = \inf \iota(\Upsilon) = \underline{x}$, by A5 there exists $\alpha' \in (0,1]$ and $h' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(\tilde{u}(A_{\alpha'}h'), \dagger_0)$, so, by the definition of σ, $\sigma(\tilde{u}(A_{\alpha'}h')) \leqslant mmc(A) = \underline{x} \leqslant x$. On the other hand, if $mmc(A) \neq \inf \iota(\Upsilon)$, by construction there exists $\alpha' \in (0,1]$ and $h' \in \mathcal{A}$ such that $\tilde{u}(A_{\alpha'}h') \in \mathcal{B}$ and so $\sigma(\tilde{u}(A_{\alpha'}h')) = \underline{x} \leqslant x$. It follows that $\leqq$ satisfies property iv, as required. This completes the proof of Lemma B.1.

□

B.1.2 Construction of confidence ranking and representation

Throughout this section, and the relevant Lemmas in Appendix C, $\leqq$ is assumed to be any relation satisfying the conditions i–v of Lemma B.1. We say that a function $D : \wp(\mathcal{A}) \to \Xi$ respects $\leqq$ if, for any $A, B \in \wp(\mathcal{A})$, $D(A) \subseteq D(B)$ iff $A \leqq B$. The following lemma is central.

Lemma B.2. *There exists a centred strict continuous confidence ranking Ξ and a cautiousness coefficient $D : \wp(\mathcal{A}) \to \Xi$ respecting $\leqq$ such that, for all $A \in \wp(\mathcal{A})$,*

$$\gamma(A, \dagger_0) = \begin{cases} \sup(A, u, D(A)) & \textit{if } \sup(A, u, D(A)) \neq \varnothing \\ \{\dagger_0\} & \textit{if} \sup(A, u, D(A)) = \varnothing \end{cases} \tag{8}$$

where u is as in Lemma C.1.

Proof. By Lemma C.1, there exists a strictly increasing zeroed continuous affine utility function u representing choices on menus consisting of constant acts. We use the notation introduced in the proof of Lemma B.1; in particular, we denote by $\tilde{u} : \wp(\mathcal{A}) \to \wp(B(K))$ the function on menus generated by u. As above, we use this function to map γ into a choice correspondence for costly deferral on $\wp(B(K)) \times \mathcal{D}$, which we also call γ. Let $\leqq$ on $\wp(B(K))$ be the image of $\leqq$ under u, and similarly for ι, mmc, Υ, Φ and $\bar{\gamma}^x$. We begin by stating two consequences of the construction and properties of $\leqq$.

Firstly, by Lemma C.4, for all $A \in \wp(B(K))$, $\alpha \in (0,1]$ and $l \in B(K)$, if $\alpha A + (1-\alpha)l \leqq A$ and $\dagger_0 \notin \gamma(A, \dagger_0)$, then $\dagger_0 \notin \gamma(\alpha A + (1-\alpha)l, \dagger_0)$, and similarly for the case where $\alpha A + (1-\alpha)l \geqq A$. Moreover, Lemma C.5 establishes that, for every $A, A' \in \Phi$, there exists $\alpha \in (0,1]$ and $l \in B(K)$ such that $\alpha A + (1-\alpha)l \equiv A'$.

Let $\mathcal{S}$ be the set of equivalence classes of $\leqq$. As standard, $\leqq$ on $\wp(B(K))$ 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 $A \in r$ and $A' \in s$, $A \leqq A'$. $r \in \mathcal{S}$ 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. By property v of $\leqq$, there is a minimal element of $\mathcal{S}$, which we call $\underline{\mathcal{S}}$; the maximal element, if it exists, is denoted by $\overline{\mathcal{S}}$. We say that an element $r \in \mathcal{S}$ is *full* if, for every $A \in \Phi$, there exists $\alpha \in (0,1]$ and $l \in B(K)$ such that $\alpha A + (1-\alpha)l \in r$. Let $\mathcal{S}^f$ be the set of full elements of $\mathcal{S}$. By Lemma C.5 and the construction of $\leqq$, every non-maximal element of $\mathcal{S}$ is full. Let $\mathcal{S}^+$ be the set of non-minimal elements of $\mathcal{S}^f$. For each $r \in \mathcal{S}^f$, define $\gamma_r : \wp(B(K)) \to 2^{B(K)}$ as follows: for all $A \in \Phi$, $\gamma_r(A) = B$ if there exists $B \in 2^{B(K)}$, $l \in B(K)$ and $\alpha \in (0,1]$ such that $\gamma(\alpha A + (1-\alpha)l, \dagger_0) = \alpha B + (1-\alpha)l$ and $\alpha A + (1-\alpha)l \in r$, and $\gamma_r(A) = \varnothing$ otherwise; and for all $A \notin \Phi$, $\gamma_r(A) = \gamma(A, \dagger_0)$. Lemma C.6 guarantees that $\gamma(\alpha A + (1-\alpha)l, \dagger_0) = \alpha B + (1-\alpha)l$ if and only if $\gamma(\beta A + (1-\beta)m, \dagger_0) = \beta B + (1-\beta)m$ for all $l, m \in B(K)$ and $\alpha, \beta \in (0,1]$ such that $\alpha A + (1-\alpha)l, \beta A + (1-\beta)m \in r$, so γ_r is well-defined for every $r \in \mathcal{S}^f$.

By Lemma C.7, the functions γ_r respect the ordering $\leqslant$ on $\mathcal{S}$ in the following sense: for all $A \in \wp(\mathcal{A})$ and $r, s \in \mathcal{S}^f$, if $r \geqslant s$, then $\gamma_r(A) \neq \varnothing$ implies that $\gamma_s(A) = \gamma_r(A)$. We now establish some further properties of the γ_r.

Lemma B.3. *For every $r \in \mathcal{S}^+$, γ_r satisfies the following continuity property: for all sequences $A_n \in \wp(B(K))$, $a_n \in B(K)$ with $A_n \to A$ and $a_n \to a$, if $a_n \in \gamma_r(A_n)$ for all $n \in \mathbb{N}$, then $a \in \gamma_r(A)$.*

Proof. Let r be a non-minimal element of $\mathcal{S}^f$, and let a_n, A_n be a pair of sequences with

$A_n \to A$, $a_n \to a$ and $a_n \in \gamma_r(A_n)$ for all $n \in \mathbb{N}$. We consider the case where $A \in \Phi$, and suppose without loss of generality that $A \in r$. (The case where $A \notin \Phi$ can be treated in an analogous fashion, replacing A in the argument below by any member of r.) If there exists $N \in \mathbb{N}$ such that $A_n \leqq A$ for all $n \geqslant N$, then, by Lemma C.7, $a_n \in \gamma(A_n, \dagger_0)$ for all $n \geqslant N$, so by A9 part i., $a \in \gamma(A, \dagger_0)$, and thus $a \in \gamma_r(A)$, as required.

Now suppose that there is no such N. Since r is a non-minimal element of $\mathcal{S}^f$, there exists $l \in B(K)$ and $\bar{\delta} \in (0,1)$ such that $\bar{\delta}A + (1-\bar{\delta})l < A$. Let $\eta = \min\{\delta \in (\bar{\delta}, 1] | \ \delta A + (1-\delta)l \geqq A\}$ (by the continuity of $\leqq$ this is a minimum). Consider any $\delta \in (\bar{\delta}, \eta)$; by the definition of η, $\delta A + (1-\delta)l < A$. By the continuity of $\leqq$, there exists $N_\delta \in \mathbb{N}$ such that, for all $n > N_\delta$, $\delta A_n + (1-\delta)l < A$. Lemma C.7 implies that $\delta a_n + (1-\delta)l \in \gamma(\delta A_n + (1-\delta)l, \dagger_0)$ for every $n > N_\delta$, whence, by A9 part i., $\delta a + (1-\delta)l \in \gamma(\delta A + (1-\delta)l, \dagger_0)$. Since this holds for all $\delta \in (\bar{\delta}, \eta)$, A9 part i. implies that $\eta a + (1-\eta)l \in \gamma(\eta A + (1-\eta)l, \dagger_0)$. Since $\eta A + (1-\eta)l \in r$, it follows that $a \in \gamma_r(A)$, as required.

□

Lemma B.4. *For each $r \in \mathcal{S}^f$, there exists a unique reflexive, transitive binary relation $\preceq_r$ representing γ_r, in the following sense: $\gamma_r(A) = \{a \in A \mid a \succeq_r b \ \ \forall b \in A\}$. Moreover, if $r = \underline{\mathcal{S}}$, then $\preceq_r$ is complete.*

Proof. We first show that for every $r \in \mathcal{S}^f$, γ_r satisfies properties α and π in Theorem 2 (see also Hill (2012)).

Consider firstly $a \in B(K)$ and $A, B \in \wp(B(K))$ such that $a \in A \subseteq B$ and $a \in \gamma_r(B)$, and let $x_r \in \Re_{\geqslant 0}$ be such that, for any $A' \in \wp(\mathcal{A})$ such that $A' \in \Upsilon$ and $A' \in r$, $\iota(A') = x_r$ (by property i of $\leqq$, x_r is well-defined). It follows from $a \in \gamma_r(B)$ that, for all $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha B + (1-\alpha)l, \dagger_{x_r})$, $\alpha a + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$. To see this, take any $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha B + (1-\alpha)l, \dagger_{x_r})$. If $\alpha B + (1-\alpha)l \in \Upsilon$, then $\iota(\alpha B + (1-\alpha)l) \leqslant r$, so $\alpha B + (1-\alpha)l \in s \leqslant r$, and since, by Lemma C.7, $\gamma_s(B) = \gamma_r(B)$, $\alpha a + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$, contradicting the assumption that $\alpha B + (1-\alpha)l \in \Upsilon$. So $\alpha B + (1-\alpha)l \notin \Upsilon$, whence it follows from $a \in \gamma_r(B)$ by the argument in Lemma C.6 that $\alpha a + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$, as required. So $a \in \bar{\gamma}^{x_r}(B)$, from which it follows, by A1, that $a \in \bar{\gamma}^{x_r}(A)$. If $A \notin \Phi$, then $\dagger_{x_r} \notin \gamma(A, \dagger_{x_r})$, so $a \in \bar{\gamma}^{x_r}(A)$ implies that $a \in \gamma(A, \dagger_0)$. Now consider the case where $A \in \Phi$, and suppose without loss of generality that $A \in r$. By the definition of ι, A4 and A9 part i., $\dagger_{x_r} \notin \gamma(A, \dagger_{x_r})$, which, given $a \in \bar{\gamma}^{x_r}(A)$, implies that $a \in \gamma(A, \dagger_0)$. So $a \in \gamma_r(A)$, and γ_r satisfies property α.

Now consider $a, b \in B(K)$ and $A, B \in \wp(B(K))$ such that $a \in A$, $b \in A \cap B$, $a \in \gamma_r(A)$ and $b \in \gamma_r(B)$. Let $x_r \in \Re_{\geqslant 0}$ be such that, for any $A' \in \wp(\mathcal{A})$ such that $A' \in \Upsilon$ and $A' \in r$, $\iota(A') = x_r$. As above, for all $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha A + (1-\alpha)l, \dagger_{x_r})$, $\alpha a + (1-\alpha)l \in \gamma(\alpha A + (1-\alpha)l, \dagger_0)$, and for all $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha B + (1-\alpha)l, \dagger_{x_r})$, $\alpha b + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$. So $a \in \bar{\gamma}^{x_r}(A)$ and $b \in \bar{\gamma}^{x_r}(B)$, whence, by A2, $a \in \bar{\gamma}^{x_r}(A \cup B)$. If $A \cup B \notin \Phi$, then $\dagger_{x_r} \notin \gamma(A \cup B, \dagger_{x_r})$, which, given $a \in \bar{\gamma}^{x_r}(A \cup B)$, implies $a \in \gamma(A \cup B, \dagger_0)$. Now consider the case where $A \cup B \in \Phi$ and suppose without loss of generality that $A \cup B \in r$. By the definition of ι, A4 and A9 part i., $\dagger_{x_r} \notin \gamma(A \cup B, \dagger_{x_r})$, which, given $a \in \bar{\gamma}^{x_r}(A \cup B)$, implies that $a \in \gamma(A \cup B, \dagger_0)$. So $a \in \gamma_r(A \cup B)$, and γ_r satisfies property π.

Now consider the case where r is non-minimal in $\mathcal{S}$. By Lemma B.3, γ_r satisfies the continuity condition in Theorem 2. By A8 part ii., for every $a \in B(K)$, $a \in \gamma(\{a\}, \dagger_0)$, so $a \in \gamma_r(\{a\})$, and γ_r satisfies $sing$. Theorem 2 implies that there exists a unique reflexive, transitive binary relation representing γ_r, as required.

For the case where $r = \underline{\mathcal{S}}$, note firstly that, by A5, for every $A \in \wp(\mathcal{A})$, there exists $\alpha \in (0,1]$, $h \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(A_\alpha h, \dagger_0)$. By Lemma C.7, it follows that $\gamma_{\underline{\mathcal{S}}}(A) \neq \varnothing$ for all $A \in \wp(\mathcal{A})$: $\gamma_{\underline{\mathcal{S}}}$ always takes non-empty values. Theorem 2 implies the required representation.

□

Note that, by A8 part i., $\preceq_r$ is non-trivial for all $r \in \mathcal{S}$. We now establish some other properties of the relations $\preceq_r$. Recall that a binary relation $\preceq$ on $B(K)$ is

- *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 $\alpha a + (1-\alpha)c \preceq \alpha b + (1-\alpha)c$.
- *continuous* if, for all $a_n, b_n \in B(K)$, if $a_n \preceq b_n$ for all $n \in \mathbb{N}$, $a_n \to a$ and $b_n \to b$, then $a \preceq b$.

Lemma B.5. *For every $r \in \mathcal{S}^f$, $\preceq_r$ is monotonic and affine. Moreover, if $r \in \mathcal{S}^+$, $\preceq_r$ is continuous.*

Proof. Monotonicity. Suppose that $a \leqslant b$. So $\alpha a + (1-\alpha)l \leqslant \alpha b + (1-\alpha)l$ for all $l \in B(K)$ and $\alpha \in (0,1]$. A8 part ii. implies that $\alpha b + (1-\alpha)l \in \gamma(\alpha\{a,b\} + (1-\alpha)l, \dagger_0)$ for all such l and α. It follows from the definition of γ_r and Lemma B.4 that $a \preceq_r b$.

Affineness. We consider the case where $\{a,b\}, \{\alpha a + (1-\alpha)c, \alpha b + (1-\alpha)c\} \in \Phi$; the other case is treated similarly. Note that it follows from the specification of the case and A8 part ii. that $a \neq b$. Since $r \in \mathcal{S}^f$, there exists $\beta \in (0,1]$ and $l \in B(K)$ such that $\beta\{a,b\} + (1-\beta)l \in r$. Consider $\{\beta(\alpha a+(1-\alpha)c)+(1-\beta)l, \beta(\alpha b+(1-\alpha)c)+(1-\beta)l\}$: since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1]$ and $m \in B(K)$ such that $\delta\{\beta(\alpha a + (1-\alpha)c + (1-\beta)l, \beta(\alpha b + (1-\alpha)c + (1-\beta)l\} + (1-\delta)m \in r$. Note that $\delta(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m = \alpha\delta(\beta a + (1-\beta)l) + (1-\alpha\delta)(\frac{\delta-\alpha\delta}{1-\alpha\delta}(\beta c + (1-\beta)l) + \frac{1-\delta}{1-\alpha\delta}m)$, where $\frac{\delta-\alpha\delta}{1-\alpha\delta}(\beta c + (1-\beta)l) + \frac{1-\delta}{1-\alpha\delta}m \in B(K)$ 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 $\frac{\delta-\alpha\delta}{1-\alpha\delta}(\beta c+(1-\beta)l) + \frac{1-\delta}{1-\alpha\delta}m = u \circ h$. Since $\{f,g\} \equiv \{f_{\alpha\delta}h, g_{\alpha\delta}h\}$, by Lemma C.4 and A3, $\beta b + (1-\beta)l \in \gamma(\beta\{a,b\} + (1-\beta)l, \dagger_0)$ iff $\delta(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m \in \gamma(\{\delta(\beta(\alpha a+(1-\alpha)c)+(1-\beta)l)+(1-\delta)m, \delta(\beta(\alpha b+(1-\alpha)c)+(1-\beta)l)+(1-\delta)m\}, \dagger_0)$. But since $\delta(\beta(\alpha a+(1-\alpha)c)+(1-\beta)l)+(1-\delta)m = \beta\delta(\alpha a+(1-\alpha)c)+(1-\beta\delta)(\frac{\delta-\beta\delta}{1-\beta\delta}l + \frac{1-\delta}{1-\beta\delta}m)$, 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.

Continuity. Continuity of γ_r for $r \in \mathcal{S}^+$ is an immediate consequence of Lemma B.3 and the representation in Lemma B.4.

□

Lemma B.6. *For each $r \in \mathcal{S}^+$, there exists a unique closed convex set of probabilities $\mathcal{C}_r$ such that, for all $A \in \wp(B(K))$,*

(9) $$\gamma_r(A) = \left\{ a \in A \mid \sum_{s \in S} a(s)p(s) \geqslant \sum_{s \in S} b(s)p(s) \quad \forall p \in \mathcal{C}_r,\ \forall b \in A \right\}$$

Proof. By A8 part i. and Lemmas B.4 and B.5, $\preceq_r$ is a non-trivial, monotonic, affine, continuous, reflexive, transitive relation, for each $r \in \mathcal{S}^+$. By Ghirardato, Maccheroni, and Marinacci (2004, Proposition A.2), for every such relation $\preceq_r$, there is a unique closed convex set of probabilities $\mathcal{C}_r$ such that, for all $a, b \in B$, $a \preceq_r b$ iff

$$(10) \qquad \sum_{s\in S} a(s)p(s) \leqslant \sum_{s\in S} b(s)p(s) \qquad \text{for all } p \in \mathcal{C}_r$$

The required representation follows from Lemma B.4.

□

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

Proof. By Lemma C.7 and the representation in Lemma B.4, for all $r, s \in \mathcal{S}^f$ with $r \geqslant s$, $\preceq_r \subseteq \preceq_s$. The result follows directly by Ghirardato, Maccheroni, and Marinacci (2004, Proposition A.1). □

Lemma B.8. *Let $\preceq_{\bigcap\mathcal{S}}$ be the relation on $B(K)$ generated by (10) with the set of probability measures $\bigcap_{r\in\mathcal{S}^+}\mathcal{C}_r$. Then $\preceq_{\underline{\mathcal{S}}} = \preceq_{\bigcap\mathcal{S}}$.*

Proof. By Lemma C.7, $\preceq_{\underline{\mathcal{S}}} \supseteq \bigcup_{r\in\mathcal{S}^+} \preceq_r$, and so $\preceq_{\underline{\mathcal{S}}} \supseteq \preceq_{\bigcap\mathcal{S}}$. For the inverse containment, suppose that $a \preceq_{\underline{\mathcal{S}}} b$. If $\{a, b\} \notin \Phi$, the result is immediate, so suppose not. Since $\preceq_{\underline{\mathcal{S}}}$ is affine, by Lemma B.5, it follows that $x^* \succeq_{\underline{\mathcal{S}}} \alpha(a-b)+x^*$, for any $\alpha \in (0, 1)$, $x \in \Re$ such that $x^*, \alpha(a-b)+x^* \in B(K)$. Take α and x such that this is the case, and let $y = \min_{s\in S}(\alpha(a-b) + x^*)(s)$. We may suppose without loss of generality that $\{x^*, \alpha(a-b) + x^*\} \in \underline{\mathcal{S}}$. Let $f, g, h \in \mathcal{A}$ such that $u \circ f = x^*$, $u \circ g = \alpha(a-b) + x^*$ and $u \circ h = y^*$. By A9 part ii., since $f \in \gamma(\{f, g\}, \dagger_0)$, $f(s) \in \gamma(\{f(s), h(s)\}, \dagger_0)$ and $g(s) \in \gamma(\{g(s), h(s)\}, \dagger_0)$ for all $s \in S$, for every $\beta \in (0, 1)$, there exists $z \in \Re_{\geqslant 0}$ such that $\bar{\gamma}^z(\{f, g\}) = \varnothing$ but $\bar{\gamma}^z(\{f, g_\beta h\}) \neq \varnothing$. It follows from the property i of $\leqq$ and its continuity that, for every $\beta \in (0, 1)$, there exists $r > \underline{\mathcal{S}}$ such that $\gamma_r(\{f, g_\beta h\}) \neq \varnothing$. By Lemma C.7, the representation in Lemma B.4 and the monotonicity and transitivity of $\preceq_{\underline{\mathcal{S}}}$, $x^* \succeq_r \beta(\alpha(a-b) + x^*) + (1-\beta)y^*$, and hence $x^* \succeq_{\bigcap\mathcal{S}} \beta(\alpha(a-b) + x^*) + (1-\beta)y^*$. Since this holds for every $\beta \in (0, 1)$ and $\preceq_{\bigcap\mathcal{S}}$ is continuous, it follows that $x^* \succeq_{\bigcap\mathcal{S}} \alpha(a-b) + x^*$ and thus, since it is affine, $a \preceq_{\bigcap\mathcal{S}} b$, as required.

□

Lemma B.9. *The set $\bigcap_{r\in\mathcal{S}^+}\mathcal{C}_r$ is a singleton.*

Proof. By Lemma B.4, $\preceq_{\underline{\mathcal{S}}}$ is complete. Since, by Lemma B.8, $\preceq_{\underline{\mathcal{S}}}$ is represented according to (10) with set of priors $\bigcap_{r\in\mathcal{S}^+}\mathcal{C}_r$, it follows that $\bigcap_{r\in\mathcal{S}^+}\mathcal{C}_r$ is a singleton.

□

Lemmas C.8 and C.9 ensure that for all $r \in \mathcal{S}^+$, $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$ and $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$ whenever r is non-maximal. Lemma C.10 establishes that, for all $r, s \in \mathcal{S}^+$, if $\mathcal{C}_r \subset \mathcal{C}_s$, then $(\mathcal{C}_r \cap (ri(\mathcal{C}_s))^c) \cap ri(\overline{\bigcup_{r' \in \mathcal{S}^f} \mathcal{C}_{r'}}) = \varnothing$, and similarly for $\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r$ and $\mathcal{C}_s$ with $s \in \mathcal{S}^+$.

Conclusion of the proof of Lemma B.2. Define

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

where the $\mathcal{C}_r$ are as specified in Lemma B.6, and where, in the second case, $\overline{\mathcal{S}}$ is understood to be a maximal element of $\mathcal{S}$ not belonging to $\mathcal{S}^f$ (ie. a non-full maximal element). It follows from Lemma B.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 B.6), Ξ is a confidence ranking. By Lemma B.9, it contains a singleton set, and so is centred. By Lemmas C.8 and C.9, Ξ is continuous, and by Lemma C.10, it is strict.

Define D as follows: for all $A \in \wp(\mathcal{A})$, if $[u \circ A] \in \mathcal{S}^+$, then $D(A) = \mathcal{C}_{[u \circ A]}$, if $A \in \underline{\mathcal{S}}$, then $D(A) = \bigcap_{s \in \mathcal{S}^+} \mathcal{C}_s$, and if $A \in \overline{\mathcal{S}}$, then $D(A) = \overline{\bigcup_{s \in \mathcal{S}^+} \mathcal{C}_s}$. By construction and Lemma B.7, D respects $\leqq$. Extensionality, Continuity and Φ-Richness of D are immediate from the definition, the fact that D respects $\leqq$, and properties iii, iv and the continuity of $\leqq$. By construction, A6, and Lemmas B.6 and B.8, u, Ξ, D represent the restriction of γ to $\wp(\mathcal{A}) \times \{\dagger_0\}$ according to (8).

□

Define $c : \Xi \to \Re_{\geqslant 0} \cup \{\infty\}$ as follows. For every $\mathcal{C} \in \Xi$, if $\mathcal{C} = \bigcap_{\mathcal{C}' \in \Xi} \mathcal{C}'$, then $c(\mathcal{C}) = u(\inf \iota(\Upsilon))$; if $\mathcal{C} = \overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}$, then $c(\mathcal{C}) = u(\sup \iota(\Upsilon))$; otherwise, $c(\mathcal{C}) = u(\iota(A))$ for any $A \in D^{-1}(\mathcal{C}) \cap \Upsilon$. (Lemma B.2, the continuity and Φ-richness of D and the fact that it respects $\leqq$, and properties i and ii of $\leqq$ imply that the value of c in the last case is well-defined: there always exists an appropriate A and the value is independent of the choice of A satisfying the conditions. The fact that u is zeroed implies that c takes non negative values.) By property i of $\leqq$ and the fact that D respects $\leqq$, c is order-preserving and -reflecting with respect to $\subseteq$. Continuity of c follows from the definition and the continuity of $\leqq$, Ξ and D.

Lemma B.10. *For every $A \in \wp(\mathcal{A})$ and $x > 0$, if $c(D(A)) \leqslant u(x)$, then*

$$\gamma(A, \dagger_x) = \begin{cases} \sup(A, u, D(A)) & \textit{if } \sup(A, u, D(A)) \neq \varnothing \\ \sup(A, u, \{p_\Xi\}) & \textit{if} \sup(A, u, D(A)) = \varnothing \end{cases}$$

where u, Ξ and D are as in Lemmas C.1 and B.2.

Proof. Let $A \in \wp(\mathcal{A})$ and $x > 0$ be such that $c(D(A)) \leqslant u(x)$. It follows from the definitions of c and ι, A4 and A9 part i. that either $\dagger_0 \notin \gamma(A, \dagger_0)$ or $\dagger_0 \in \gamma(A, \dagger_0)$ and $\dagger_x \notin \gamma(A, \dagger_x)$. It follows from Lemma B.2 that $\sup(A, u, D(A)) \neq \varnothing$ in the former case and that $\sup(A, u, D(A)) = \varnothing$ in the latter case. Moreover, in the former case, A4 implies that $\gamma(A, \dagger_{x'}) = \sup(A, u, D(A))$, for all $x' \geqslant 0$, which implies the desired representation. Now consider the latter case. Note that, in this case, $A \in \Phi$.

By the Φ-Richness of D, there exists $\alpha \in (0, 1]$ and $h \in \mathcal{A}$ such that $D(A_\alpha h) = \bigcap_{\mathcal{C}' \in \Xi} \mathcal{C}'$. By the fact that Ξ is centred, $\bigcap_{\mathcal{C}' \in \Xi} \mathcal{C}' = \{p_\Xi\}$. Lemma B.2 implies that $\gamma(A_\alpha h, \dagger_0) = \sup(A_\alpha h, u, \{p_\Xi\}))$. By A3, since $\dagger_x \notin \gamma(A, \dagger_x)$, we have that $\gamma(A, \dagger_x)_\alpha h = \gamma(A_\alpha h, \dagger_0) = \sup(A_\alpha h, u, \{p_\Xi\})$. So $\gamma(A, \dagger_x) = \sup(A, u, \{p_\Xi\})$, as required.

□

The desired representation is a direct consequence of Lemmas B.2 and B.10, A6, and the construction of c. Hence the axioms imply the representation, as required.

B.2 Necessity of axioms

As concerns the (ii) to (i) direction, most cases are straightforward, in the light of the fact that the representation (2) implies that, if $f \in \bar{\gamma}^x(A, \dagger_0)$, then f is an optimal element in A according to the unanimity rule with the level of confidence $c^{-1}(x)$. So A1, for example, holds if $f \in \sup(B, u, \mathcal{C})$ implies that $f \in \sup(A, u, \mathcal{C})$, where $\mathcal{C}$ is the confidence level such that $c(\mathcal{C}) = u(x)$; this is evidently the case. Similar considerations apply to A2 and A9 part ii. (in the latter case, note that the fact above implies that $mmc(A)$ gives the cost of the highest confidence level which yields an optimal element from A according to the unanimity rule). Apart from the continuity axioms, which shall be discussed below, perhaps the only other axiom requiring comment is A3. Note that, since the confidence ranking is strict, if $\sum_{s \in S} u(f(s)).p(s) = u(c)$ for all $p \in \mathcal{C}$, for some $\mathcal{C} \in \Xi$ and $f \in \mathcal{A}$, $c \in \Delta(X)$, then for any $\mathcal{C}' \in \Xi$ such that $\sum_{s \in S} u(f(s)).p(s) \neq u(c)$ for some $p \in \mathcal{C}'$, there exists $q, q' \in \mathcal{C}'$ such that $\sum_{s \in S} u(f(s)).q(s) > u(c)$ and $\sum_{s \in S} u(f(s)).q'(s) < u(c)$. Hence, if $\sup(\{f, c\}, u, \mathcal{C}) = \{f, c\}$ for some $\mathcal{C} \in \Xi$, then for every $\mathcal{C}' \in \Xi$, either $\sup(\{f, c\}, u, \mathcal{C}') =$

$\{f, c\}$ or $\sup(\{f, c\}, u, \mathcal{C}') = \varnothing$. It follows, given the properties of the unanimity rule, that for any $A \in \wp(\mathcal{A})$, if $\sup(A, u, \mathcal{C}) = A'$ for some $\mathcal{C} \in \Xi$, then for every $\mathcal{C}' \in \Xi$, either $\sup(A, u, \mathcal{C}') = A'$ or $\sup(A, u, \mathcal{C}') = \varnothing$. In the light of representation (2), this implies that A3 holds.

The final cases of potential interest are A9, parts i and iii.[33] As concerns the former, consider any $A_n, A \in \wp(\mathcal{A})$, $f_n, f \in \mathcal{A}$ and $x_n, x \in \Re_{\geqslant 0}$, $n \in \mathbb{N}$ with $A_n \to A$, $f_n \to f$ and $x_n \to x$, and such that $f_n \in \gamma(A_n, \dagger_{x_n})$ for all $n \in \mathbb{N}$. We distinguish two cases. If $\dagger_0 \in \gamma(A_n, \dagger_0)$ for all $n \geqslant N$ for some $N \in \mathbb{N}$, then $f \in \gamma(A, \dagger_x)$ follows from the continuity of the EU representation, the continuity of c, the form of representation (2), and (in the case where $\dagger_0 \notin \gamma(A, \dagger_0)$) the fact, noted above, that $\sup(A, u, \{p_\Xi\}) = \sup(A, u, \mathcal{C})$ for any $\mathcal{C} \in \Xi$ such that $\sup(A, u, \mathcal{C}) \neq \varnothing$. Now consider the case where, for every $N \in \mathbb{N}$, there exists $n > N$ with $\dagger_0 \notin \gamma(A_n, \dagger_0)$. Given the form of representation (2), it suffices to show that, for any sequences $A_n \to A$ and $f_n \to f$ with $f_n \in \gamma(A_n, \dagger_0)$ for all $n \in \mathbb{N}$, $f \in \gamma(A, \dagger_0)$. Taking such a pair of sequences, we consider the case where there is no $N \in \mathbb{N}$ with $D(A_n) \supseteq D(A)$ for all $n \geqslant N$; the other case is treated similarly. Since there is no such sequence, $D(A) \neq \{p_\Xi\}$, so there exists $B \in \wp(\mathcal{A})$ with $D(B) \subset D(A)$. By the continuity of D, for each such B, there exists $N_B \in \mathbb{N}$ such that $D(A_n) \supseteq D(B)$ for all $n \geqslant N_B$. It follows that, for each $n \geqslant N_B$, $\sum_{s \in S} u(f_n(s)).p(s) \geqslant \sum_{s \in S} u(h(s)).p(s)$ for all $p \in D(B)$ and $h \in A_n$. Hence, by the continuity of the representation, $\sum_{s \in S} u(f(s)).p(s) \geqslant \sum_{s \in S} u(h(s)).p(s)$ for all $p \in D(B)$ and $h \in A$. Since this holds for every B with $D(B) \subset D(A)$, and since the continuity of the confidence ranking and D imply that $D(A) = \overline{\bigcup_{D(B) \subset D(A)} D(B)}$, there cannot be a $q \in D(A)$ such that $\sum_{s \in S} u(f(s)).q(s) < \sum_{s \in S} u(h(s)).q(s)$ for some $h \in A$. So $f \in \sup(A, u, D(A))$; hence $f \in \gamma(A, \dagger_0)$ and so $f \in \gamma(A, \dagger_x)$, as required.

Finally, we consider A9 part iii. We use the fact, established in Lemma C.11, that $mmc^{-1}([0, x])$ and $mmc^{-1}([x, \infty))$ are closed in Φ for all $x \in (0, \sup \iota(\Upsilon))$. Suppose that A9 part iii. does not hold, and that there exists $A \in \Phi$ with $A \in \overline{(\iota^{-1}([0, x]) \cap \Upsilon) \cup (mmc^{-1}([0, x]) \cap \Upsilon^c)} \backslash (\iota^{-1}([0, x]) \cap \Upsilon) \cup (mmc^{-1}([0, x]) \cap \Upsilon^c)$ for some $x \in (0, \sup \iota(\Upsilon))$. It follows that $A \in \overline{(\iota^{-1}([0, x]) \cap \Upsilon)} \backslash (\iota^{-1}([0, x]) \cap \Upsilon)$ or $A \in \overline{(mmc^{-1}([0, x]) \cap \Upsilon^c)} \backslash (mmc^{-1}([0, x]) \cap \Upsilon^c)$. Suppose first that $A \in$

[33] Richness (A10) is a straightforward consequence of the Φ-Richness of D for part i., and of the representation and the Φ-Richness and continuity of D for part ii. Likewise, A9 part iv. is a straightforward consequence of the Φ-Richness of D and the form of the representation.

$\overline{(\iota^{-1}([0,x]) \cap \Upsilon)} \backslash (\iota^{-1}([0,x]) \cap \Upsilon)$. If $A \in \Upsilon$, then, by the continuity of D and c and the representation, $A \in \iota^{-1}([0,x]) \cap \Upsilon$. If $A \notin \Upsilon$, then, since $\iota^{-1}([0,x]) \subseteq mmc^{-1}([0,x])$ and the latter set is closed by Lemma C.11, $A \in mmc^{-1}([0,x]) \cap \Upsilon^c$. Now suppose that $A \in \overline{(mmc^{-1}([0,x]) \cap \Upsilon^c)} \backslash (mmc^{-1}([0,x]) \cap \Upsilon^c)$. Since, by the continuity of the representation, Υ^c is closed, $A \in \Upsilon^c$. It follows, since $mmc^{-1}([0,x])$ is closed (by Lemma C.11), that $A \in mmc^{-1}([0,x]) \cap \Upsilon^c$. All cases contradict the assumption that $A \in \overline{(\iota^{-1}([0,x]) \cap \Upsilon) \cup (mmc^{-1}([0,x]) \cap \Upsilon^c)} \backslash (\iota^{-1}([0,x]) \cap \Upsilon) \cup (mmc^{-1}([0,x]) \cap \Upsilon^c)$; so $(\iota^{-1}([0,x]) \cap \Upsilon) \cup (mmc^{-1}([0,x]) \cap \Upsilon^c)$ is closed in Φ, as required. A similar argument establishes that $(\iota^{-1}([x,\infty]) \cap \Upsilon) \cup (mmc^{-1}([x,\infty]) \cap \Upsilon^c)$ is closed in Φ.

B.3 Uniqueness

Uniqueness of u follows from the Herstein-Milnor theorem. Without loss of generality, we can thus restrict attention to two representations (u, Ξ, D, c) and (u, Ξ', D', c') of the same choice correspondence γ involving the same utility function. Note firstly that if $\Phi = \varnothing$, then uniqueness of all elements follows from the standard uniqueness properties of the EU representation; henceforth we suppose that $\Phi \neq \varnothing$. (2) implies that, for any $A \in \Upsilon$, $c(D(A)) = u(\iota(A))$. So the function $c \circ D : \wp(\mathcal{A}) \to \Re_{\geqslant 0} \cup \{\infty\}$ is unique on Υ. By the continuity of c and D, it follows that it is unique on $\overline{\Upsilon}$.

We now consider the uniqueness of D. We first establish that D and D' coincide on Υ. Suppose, for the purposes of reductio, that there exists $\hat{A} \in \Upsilon$ for which they do not coincide. Suppose, without loss of generality, that $p \in D(\hat{A}) \backslash D'(\hat{A})$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in D'(\hat{A})$. Since B is finite-dimensional, 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 and α can be chosen so that $\alpha \in K$ and $a \in B(K)$. Taking $f \in \mathcal{A}$ such that $u \circ f = a$ and $c \in \Delta(X)$ such that $u(c) = \alpha$, we have that $\sum_{s \in S} u(f(s))p(s) \geqslant u(c)$ for all $p \in D'(\hat{A})$, whereas this is not the case for all $p \in D(\hat{A})$. Let $x \in \Re_{>0}$ be such that $u(x) = c(D(\hat{A})) = c'(D'(\hat{A}))$ (note that since $\hat{A} \in \Upsilon$, it belongs to the domain where $c \circ D$ and $c' \circ D'$ agree). It follows from the representation (2) and the aforementioned properties of $D'(\hat{A})$ that $f \in \bar{\gamma}^x(\{f, c\})$. However, it follows from the representation and the properties of $D(\hat{A})$ that $f \notin \bar{\gamma}^x(\{f, c\})$, contradicting the assumption that (u, Ξ, D, c) and (u, Ξ', D', c') both represent the choice correspondence γ. Hence D and D' coincide on Υ. Now consider $A \in \overline{\Upsilon} \backslash \Upsilon$, and take any sequence of $A_n \in \Upsilon$ with $A_n \to A$. By the continuity of D, $D(A) = \lim_{n\to\infty} D(A_n)$ and similarly for D'; since D and D' coincide on Υ, $D(A) = D'(A)$. So D and D' coincide on $\overline{\Upsilon}$.

As concerns the uniqueness of Ξ, we show that $\Xi = \Xi' = D(\overline{\Upsilon})$. It follows from what has just been established that Ξ and Ξ' both contain $D(\overline{\Upsilon})$. We firstly show that there do not exist $\mathcal{C} \in \Xi \backslash D(\overline{\Upsilon})$ with $\mathcal{C} \subsetneq \bigcap_{A' \in \overline{\Upsilon}} D(A')$. Suppose that there does exist such $\mathcal{C} \in \Xi$. Using a separating hyperplane theorem as above, one can construct $f \in \mathcal{A}$, $c \in \Delta(X)$ such that $\sum_{s\in S} u(f(s))p(s) \geqslant u(c)$ for all $p \in \mathcal{C}$, whereas this is not the case for $\bigcap_{A' \in \overline{\Upsilon}} D(A')$. Hence $\{f, c\} \in \Phi$. By the continuity of Ξ, there exists $\mathcal{C}'' \in \Xi$ such that $\mathcal{C}'' \subsetneq \bigcap_{A' \in \overline{\Upsilon}} D(A')$ and it is not the case that $\sum_{s\in S} u(f(s))p(s) \geqslant u(c)$ for all $p \in \mathcal{C}''$. By the Φ-Richness and continuity of D, there exists $\alpha \in (0, 1]$, $h \in \mathcal{A}$ such that $D(\{f_\alpha h, c_\alpha h\}) = \mathcal{C}''$. It follows from the representation that $\{f_\alpha h, c_\alpha h\} \in \Upsilon$, and from the order-preserving and -reflecting properties of c that $\iota(\{f_\alpha h, c_\alpha h\}) < \inf \iota(\Upsilon)$, which is a contradiction. So there is no $\mathcal{C} \in \Xi \backslash D(\overline{\Upsilon})$ with $\mathcal{C} \subsetneq \bigcap_{A' \in \overline{\Upsilon}} D(A')$, as required. A similar argument shows that there exists no $\mathcal{C} \in \Xi \backslash D(\overline{\Upsilon})$ with $\mathcal{C} \supsetneq \bigcup_{A' \in \overline{\Upsilon}} D(A')$. Since Ξ is nested, it follows that $\Xi = D(\overline{\Upsilon})$, and similarly for Ξ'. So $\Xi = \Xi'$, as required.

Finally, the uniqueness of c is a direct consequence of the fact that $\Xi = D(\overline{\Upsilon})$, and the uniqueness of D and $c \circ D$ on $\overline{\Upsilon}$.

□

Appendix C Proofs of results mentioned in Appendices A and B

Throughout this Appendix, we adopt the notation introduced in Appendix B.

C.1 Proofs of results used in Appendix B

Proof of Theorem 2. Define $\preceq$ by $x \preceq y$ iff $y \in \gamma(\{x, y\})$. Reflexivity is an immediate consequence of $sing$.[34] Transitivity follows from π and α, by the same reasoning as used

[34] Note that there is an error in the statement and proof of Theorem 2 in Hill (2012). The proof (p300) establishes the representation for all non-singleton menus, but not for singleton menus; accordingly, it does not show the reflexivity of the representing relation. This omission is corrected by the addition of the axiom $sing$.

in the proof of Hill (2012, Theorem 2). It is a straightforward consequence of α that, if $x \in \gamma(A)$, then $x \succeq y$ for all $y \in A$. It remains to show that if $x \succeq y$ for all $y \in A$, then $x \in \gamma(A)$.

Suppose that this is not the case: ie. $x \succeq y$ for all $y \in A$, but $x \notin \gamma(A)$. We first show that there exists a maximal subset A' of A such that $x \in A'$ and $x \in \gamma(A')$. $sing$ implies that there exists at least one subset such that x is in the image of γ, namely $\{x\}$. The continuity of γ implies that if there were an increasing (under set inclusion) infinite chain A_i of subsets of A such that $x \in \gamma(A_i)$ for all i, then $x \in \gamma(\overline{\bigcup A_i})$ (where $\overline{\bigcup A_i}$ is compact since it is contained in A) so this is a subset containing all the subsets in the chain and such that x is in its image under γ. Hence, by Zorn's Lemma, there exists a maximal subset of A such that $x \in A'$ and $x \in \gamma(A')$. Since $x \notin \gamma(A)$, any such maximal subset cannot be A. Take any such subset A' and consider any $y \in A \backslash A'$. Since $x \in \gamma(A)$ and $x \in \gamma(\{x, y\})$ by hypothesis, π implies that $x \in \gamma(A' \cup \{y\})$, contradicting the maximality of A'. Hence $x \in \gamma(A)$ as required.

The uniqueness is immediate from the definition of $\preceq$. Finally, since π implies Sen's β, the standard theorem for choice correspondences (for example Sen (1971)) implies that, if γ always takes non-empty values then it can be represented by a complete, reflexive, transitive binary relation, without the need for the continuity assumption. □

Lemma C.1. *There exists a strictly increasing zeroed continuous affine utility function $u : \Delta(X) \to \Re$ representing the restriction of γ to sets of constant acts.*

Proof. A6 and A7 imply that $\gamma(A, \dagger_0) \subseteq A$ for all $A \in \wp(\mathcal{A})$ such that $A \subseteq \Delta(X)$. A1 and A2 imply that the restriction of $\gamma(\bullet, \dagger_0)$ to sets of constant acts satisfies properties α and π in Theorem 2. By Theorem 2, there exists a complete, transitive, reflexive preference relation $\preceq^0_{|\Delta(X)}$ representing the restriction to $\gamma(\bullet, \dagger_0)$ to sets of constant acts. By A8 part i., A3 and A9 part i., $\preceq^0_{|\Delta(X)}$ is a non-trivial relation satisfying independence and continuity. The existence of an affine u representing $\preceq^0_{|\Delta(X)}$ follows from the Herstein-Milnor theorem. A9 part i. implies that u is continuous. By A8 part i., u is strictly increasing. Since u is unique up to positive affine transformation, it can be chosen to be zeroed, as required. By A4, $\gamma(A, \dagger_x) = \gamma(A, \dagger_0)$ for all $A \in \wp(\mathcal{A})$ with $A \subseteq \Delta(X)$ and all $x \in \Re_{\geqslant 0}$, so u represents the restriction of γ to constant acts, as required.

□

Lemma C.2. *For all $A, B \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$, if $A \overset{e.e.}{\simeq} B$ with relevant correspondence σ, then $\gamma(B, \dagger_x) = \sigma(\gamma(A, \dagger_x))$ if $\gamma(A, \dagger_x) \subseteq A$ and $\gamma(B, \dagger_x) = \dagger_x$ if not.*

Proof. Take any A, B and σ with the specified properties. If $\dagger_x \in \gamma(A, \dagger_x)$, then A8 part iii. and A6 imply that $\gamma(B, \dagger_x) = \dagger_x$. Now consider the case where $\dagger_x \notin \gamma(A, \dagger_x)$. Let $\bar{A} = \gamma(A, \dagger_x)$; by A6, $\bar{A} \subseteq \mathcal{A}$. By A5, there exists $\alpha \in (0, 1]$ and $h \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(A_\alpha h, \dagger_0)$, and by A3, $\gamma(A_\alpha h, \dagger_0) = \bar{A}_\alpha h$ for any such α and h. For any $f \in \bar{A}$ and $f' \in \sigma(f)$, it follows from A5 that there exists $\alpha' \in (0, 1]$ and $h' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$. For any such α' and h', A2 (applied with $x = 0$) implies that $\bar{A}_{\alpha'} h' \cup \{f'_{\alpha'} h'\} \subseteq \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$, and A1 implies that $\gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0) = \bar{A}_{\alpha'} h' \cup \{f'_{\alpha'} h'\}$. Similarly, for any $f \in A \backslash \bar{A}$ and $f' \in \sigma(f)$, A5 implies that there exists $\alpha' \in (0, 1]$ and $h' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$. For any such α' and h', A2 implies that $\bar{A}_{\alpha'} h' \subseteq \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$ and that if $f'_{\alpha'} h' \in \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$, then $f_{\alpha'} h' \in \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$, from which it would follow by A1 that $f_{\alpha'} h' \in \bar{A}_{\alpha'} h'$, contrary to the assumption. Hence $\gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0) = \bar{A}_{\alpha'} h'$. Repeating this reasoning, we have that, for $\alpha'' \in (0, 1]$ and $h'' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma((A \cup B)_{\alpha''} h'', \dagger_0)$, $\gamma((A \cup B)_{\alpha''} h'', \dagger_0) = \bar{A}_{\alpha''} h'' \cup \sigma(\bar{A})_{\alpha''} h''$. Hence, by A1, for $\alpha''' \in (0, 1]$ and $h''' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(B_{\alpha'''} h''', \dagger_0)$, $\gamma(B_{\alpha'''} h''', \dagger_0) = \sigma(\bar{A})_{\alpha'''} h'''$. Since, by A8 part iii., $\dagger_x \notin \gamma(B, \dagger_x)$, it follows from A3 that $\gamma(B, \dagger_x) = \sigma(\bar{A})$, as required.

□

Lemma C.3. *For all $A \in \Upsilon$, $\iota(A) > mmc(A)$.*

Proof. Note that A9 part i., A4 and A6 imply that, if $\iota(A) = y$, then $\dagger_y \notin \gamma(A, \dagger_y)$. Hence, for all $A \in \wp(\mathcal{A})$, $\iota(\{A_\alpha h \mid \alpha \in (0, 1], h \in \mathcal{A}, A_\alpha h \in \Upsilon\}) \subseteq \{x \in \Re_{\geqslant 0} \mid \bar{\gamma}^x(A) = \varnothing\}$. By A9 part iv., this set is open, so, for any $A \in \Upsilon$, $\iota(A) > mmc(A)$, as required.

□

Lemma C.4. *For all $B \in \wp(\mathcal{A})$, $h \in \mathcal{A}$ and $\alpha \in (0, 1]$, if $B_\alpha h \leqq B$ and $\dagger_0 \notin \gamma(B, \dagger_0)$, then $\dagger_0 \notin \gamma(B_\alpha h, \dagger_0)$. Similarly, if $B_\alpha h \geqq B$ and $\dagger_0 \notin \gamma(B_\alpha h, \dagger_0)$, then $\dagger_0 \notin \gamma(B, \dagger_0)$.*

Proof. First note that, for any $B \in \mathcal{A}$, $h \in \mathcal{A}$ and $\alpha \in (0, 1]$, $mmc(B_\alpha h) = mmc(B)$ by A9 part iv. Now suppose that $B \in \mathcal{A}$, $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ are such that $B_\alpha h \leqq B$ and $\dagger_0 \notin \gamma(B, \dagger_0)$. We first show that $\iota(B_\alpha h) = 0$. If this were not the case, then Lemma C.3 implies that $\iota(B_\alpha h) > mmc(B_\alpha h) = mmc(B)$. It follows, by property ii of $\leqq$ (Lemma

B.1) that $B_\alpha h > B$, contradicting the fact that $B_\alpha h \leqq B$. So $\iota(B_\alpha h) = 0$. By the definition of ι and A4, there exists $B' \subseteq B$ such that $\gamma(B_\alpha h, \dagger_x) = B'_\alpha h$ for all $x > 0$, whence, by A9 part i., $\dagger_0 \notin \gamma(B_\alpha h, \dagger_0)$, as required. Similar reasoning establishes the conclusion for the other case.

□

Lemma C.5. *For every $A, A' \in \wp(B(K))$ such that $A, A' \in \Phi$, there exists $\alpha \in (0,1]$ and $l \in B(K)$ such that $\alpha A + (1-\alpha)l \equiv A'$.*

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

□

Lemma C.6. *For every $A \in \wp(B(K))$, $l, m \in B(K)$ and $\alpha, \beta \in (0,1]$ with $\alpha A + (1-\alpha)l, \beta A + (1-\beta)m \in r \in \mathcal{S}^f$, and every $a \in A$, $\alpha a + (1-\alpha)l \in \gamma(\alpha A + (1-\alpha)l, \dagger_0)$ iff $\beta a + (1-\beta)m \in \gamma(\beta A + (1-\beta)m, \dagger_0)$.*

Proof. The result is an immediate consequence of A3 if $A \notin \Phi$, so suppose that this is not the case. 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. Let $B \in \wp(\mathcal{A})$ be such that $\alpha A + (1-\alpha)l = \tilde{u}(B)$ and $h \in \mathcal{A}$ be such that $\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m = u \circ h$; so $\beta A + (1-\beta)m = \tilde{u}(B_{\frac{\beta}{\alpha}}h)$. Since $B \equiv B_{\frac{\beta}{\alpha}}h$, Lemma C.4 implies that $\dagger_0 \notin \gamma(B, \dagger_0)$ iff $\dagger_0 \notin \gamma(B_{\frac{\beta}{\alpha}}h, \dagger_0)$. Hence, by A3 and A6, $\gamma(B, \dagger_0)_{\frac{\beta}{\alpha}}h = \gamma(B_{\frac{\beta}{\alpha}}h, \dagger_0)$, which yields the required conclusion.

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 $\epsilon A + (1-\epsilon)k \equiv \alpha A + (1-\alpha)l$, then, by applying the reasoning in the case above, we get $\alpha a + (1-\alpha)l \in \gamma(\alpha A + (1-\alpha)l, \dagger_0)$ iff $\epsilon a + (1-\epsilon)k \in \gamma(\epsilon A + (1-\epsilon)k, \dagger_0)$ iff $\beta a + (1-\beta)m \in \gamma(\beta A + (1-\beta)m, \dagger_0)$, as required. By the continuity and non-triviality of $\leqq$, $\mathcal{S}^f$ is not a singleton, so there exists $n \in B(K)$ and $\eta \in (0,1)$ such that $\eta(\alpha A + (1-\alpha)l) + (1-\eta)n \not\equiv \alpha A + (1-\alpha)l$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1)$ and $n' \in B(K)$

such that $\delta(\eta(\alpha A + (1-\alpha)l) + (1-\eta)n) + (1-\delta)n' \equiv \alpha A + (1-\alpha)l \in r$. However, $\delta(\eta(\alpha A+(1-\alpha)l)+(1-\eta)n)+(1-\delta)n' = \alpha\eta\delta A+(1-\alpha\eta\delta)(\frac{\delta-\alpha\eta\delta}{1-\alpha\eta\delta}(\frac{\eta-\alpha\eta}{1-\alpha\eta}l+\frac{1-\eta}{1-\alpha\eta}n)+\frac{1-\delta}{1-\alpha\eta\delta}n')$, with $\frac{\delta-\alpha\eta\delta}{1-\alpha\eta\delta}(\frac{\eta-\alpha\eta}{1-\alpha\eta}l + \frac{1-\eta}{1-\alpha\eta}n) + \frac{1-\delta}{1-\alpha\eta\delta}n' \in B(K)$ since it is a mix of elements of $B(K)$. So $\alpha\eta\delta$ and $\frac{\delta-\alpha\eta\delta}{1-\alpha\eta\delta}(\frac{\eta-\alpha\eta}{1-\alpha\eta}l + \frac{1-\eta}{1-\alpha\eta}n) + \frac{1-\delta}{1-\alpha\eta\delta}n'$ have the properties required above, and the result is established.

□

Lemma C.7. *For all $A \in \wp(\mathcal{A})$ and $r, s \in \mathcal{S}^f$ with $r \geqslant s$, if $\gamma_r(A) \neq \varnothing$ then $\gamma_s(A) = \gamma_r(A)$.*

Proof. If $s = r$, there is nothing to show, so suppose not. If $A \notin \Phi$, the result is an immediate consequence of the definition of γ_r, so suppose this is not the case. Without loss of generality, it can be assumed that $A \in r$. (If not, apply the argument below to an $\alpha A + (1-\alpha)l \in r$.) Let $\beta \in (0,1)$ and $m \in B(K)$, be such that $\beta A + (1-\beta)m \in s$ (such β and m exist since $s \in \mathcal{S}^f$). Since, by assumption and A6, $\dagger_0 \notin \gamma(A, \dagger_0)$, Lemma C.4 implies that $\dagger_0 \notin \gamma(\beta A + (1-\beta)m, \dagger_0)$; A3 implies that $\gamma(\beta A + (1-\beta)m, \dagger_0) = \beta\gamma(A, \dagger_0) + (1-\beta)m$, so $\gamma_s(A) = \gamma_r(A)$, as required. □

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

Proof. By Lemma B.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, there is a 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 B is finite-dimensional, 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 $\alpha \in K$, $a \in B(K)$. By construction, $\{a, \alpha^*\} \in \Phi$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1]$ and $m \in B(K)$ such that $\delta\{a, \alpha^*\} + (1-\delta)m \in r$. Since r is non-minimal in $\mathcal{S}^f$, there exists $l \in B(K)$ and $\beta \in (0,1)$ such that $\beta(\delta\{a, \alpha^*\} + (1-\delta)m) + (1-\beta)l < \delta\{a, \alpha^*\} + (1-\delta)m$. Let $\beta' = \min\{\epsilon \in [\beta, 1] |\ \epsilon(\delta\{a, \alpha^*\} + (1-\delta)m) + (1-\epsilon)l \geqq \delta\{a, \alpha^*\} + (1-\delta)m\}$ (this is a minimum by the continuity of $\leqq$). Taking $f, g, h \in \mathcal{A}$ such that $u \circ f = \delta a + (1-\delta)m$, $u \circ g = \delta\alpha^* + (1-\delta)m$ and $u \circ h = l$, it follows, by the construction, that for any $\eta \in (\beta, \beta')$, $f_\eta h \in \gamma(\{f_\eta h, g_\eta h\}, \dagger_0)$. However, by construction, $f_{\beta'}h \notin \gamma(\{f_{\beta'}h, g_{\beta'}h\}, \dagger_0)$, contradicting A9 part i. Hence $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. □

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

Proof. By Lemma B.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, there is a 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 $q \in \mathcal{C}_r$. Since B is finite-dimensional, 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 $\alpha \in K$, $a \in B(K)$. By construction, $\{a, \alpha^*\} \in \Phi$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1]$ and $m \in B(K)$ such that $\delta\{a, \alpha^*\} + (1-\delta)m \in r$. Take any $x \in K$ with $x \leqslant \alpha, a(s)$ for all $s \in S$, and let $f, g, h \in \mathcal{A}$ be such that $u \circ f = \delta a + (1-\delta)m$, $u \circ g = \delta\alpha^* + (1-\delta)m$, $u \circ h = \delta x^* + (1-\delta)m$. Let $\beta \in (0,1)$ be such that $u \circ g_\beta h = \delta(\alpha - \frac{\epsilon}{2})^* + (1-\delta)m$; such a β exists by the choice of a and α and the definition of g and h. By construction, $f \in \gamma(\{f,g\}, \dagger_0)$ and for all $\alpha \in (0,1]$ and $e \in \mathcal{A}$ such that $\{f, g_\beta h\}_\alpha e > \{f,g\}$, $f_\alpha e \notin \gamma(\{f, g_\beta h\}_\alpha e, \dagger_0)$. However, $f(s) \in \gamma(\{f(s), h(s)\}, \dagger_0)$ and $g(s) \in \gamma(\{g(s), h(s)\}, \dagger_0)$ for all $s \in S$, and there exists $\alpha \in (0,1]$ and $e \in \mathcal{A}$ with $\{f, g_\beta h\}_\alpha e > \{f,g\}$. Since, by A1, A2, A3 and A8 part ii., for all $\alpha \in (0,1]$ and $e \in \mathcal{A}$, if $\dagger_0 \notin \gamma(\{f, g_\beta h\}_\alpha e, \dagger_0)$ then $f_\alpha e \in \gamma(\{f, g_\beta h\}_\alpha e, \dagger_0)$, it follows that, for all $\alpha \in (0,1]$ and $e \in \mathcal{A}$ such that $\{f, g_\beta h\}_\alpha e > \{f,g\}$, $\{f, g_\beta h\}_\alpha e \in \Upsilon$. It follows by properties i, ii and iv of $\leqq$ and its continuity that for all $\sup(\iota(\wp(\mathcal{A}))) > x > mmc(\{f,g\})$, there exists $\alpha \in (0,1]$, $e \in \mathcal{A}$ with $\iota(\{f, g_\beta h\}_\alpha e) = x$. So for every $x \in \Re_{\geqslant 0}$ such that $\bar{\gamma}^x(\{f,g\}) = \varnothing$, $\bar{\gamma}^x(\{f, g_\beta h\}) = \varnothing$, contradicting A9 part ii. Hence $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.

□

Lemma C.10. *For all $r, s \in \mathcal{S}^+$, if $\mathcal{C}_r \subset \mathcal{C}_s$, then $(\mathcal{C}_r \cap (ri(\mathcal{C}_s))^c) \cap ri(\overline{\bigcup_{r' \in \mathcal{S}^+} \mathcal{C}_{r'}}) = \varnothing$. Similarly, for all $s \in \mathcal{S}^+$, $(\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r \cap (ri(\mathcal{C}_s))^c) \cap ri(\overline{\bigcup_{r' \in \mathcal{S}^+} \mathcal{C}_{r'}}) = \varnothing$.*

Proof. We only consider the case of $r, s \in \mathcal{S}^+$, the other case being treated similarly. Suppose that the condition does not hold, so there exist $r, s \in \mathcal{S}^+$ with $\mathcal{C}_r \subset \mathcal{C}_s$ and $p \in \mathcal{C}_r \cap (ri(\mathcal{C}_s))^c$ but $p \in ri(\overline{\bigcup_{r' \in \mathcal{S}^f} \mathcal{C}_{r'}})$. Let $x_r = \iota(A)$ for any $A \in \Upsilon$ such that $A \in r$ and similarly for x_s; property i of $\leqq$ implies that these are well-defined and Lemma B.7 implies that $x_s > x_r$. Since S is finite, it follows from a supporting hyperplane theorem (Aliprantis and Border, 2007, Theorem 7.36) that there exists a linear functional ϕ supporting $\mathcal{C}_s$ at p; ie. such that $\phi(q) \geqslant \phi(p)$ for all $q \in \mathcal{C}_s$. Let $\phi(p) = \alpha$. Since B is finite-dimensional, 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 can be chosen so that $a \in B(K)$ and $\alpha \in K$. By construction, $\{a, \alpha^*\} \in \Phi$. Since $r \in \mathcal{S}^f$, there exist $\delta \in (0,1]$ and $m \in B(K)$ such that $\delta\{a, \alpha^*\} + (1-\delta)m \in r$.

Take $f, g \in \mathcal{A}$ such that $u \circ f = \delta a + (1-\delta)m$ and $u(g) = \delta\alpha^* + (1-\delta)m$. Since, by construction, $\delta a + (1-\delta)m \in \gamma_s(\delta\{a, \alpha^*\} + (1-\delta)m)$, it follows that $mmc(\{f,g\}) \geqslant x_s > x_r$. Finally, it follows from the construction that for any $z > \alpha$ and $g' \in \mathcal{A}$ such that $u(g') = \delta z^* + (1-\delta)m$, and any $\beta \in (0,1)$, $\dagger_0 \in \gamma(\{f_\alpha h, (g_\beta g')_\alpha h\}, \dagger_0)$ for all $\alpha \in (0,1]$, $h \in \mathcal{A}$ such that $\{f_\alpha h, (g_\beta g')_\alpha h\} \equiv \{f,g\}$. Since, whenever $\{f_\alpha h, (g_\beta g')_\alpha h\} \equiv \{f,g\}$, $\dagger_{x_r} \notin \gamma(\{f_\alpha h, (g_\beta g')_\alpha h\}, \dagger_{x_r})$, it follows that each $\{f_\alpha h, (g_\beta g')_\alpha h\}$ with $\{f_\alpha h, (g_\beta g')_\alpha h\} \equiv \{f,g\}$ belongs to $\iota^{-1}([0,x_r]) \cap \Upsilon$. By the continuity of $\leqq$, every neighborhood of $\{f,g\}$ contains such a $\{f_\alpha h, (g_\beta g')_\alpha h\}$; so $\{f,g\}$ is an element of Φ on the boundary of $(\iota^{-1}([0,x_r]) \cap \Upsilon) \cup (mmc^{-1}([0,x_r] \cap \Upsilon^c)$, but not belonging to this set, contradicting A9 part iii. So there exist no such r, s, as required.

□

Lemma C.11. *Let γ be represented according to (2). Then, for all $x \in (0, \sup\iota(\Upsilon))$, $mmc^{-1}([0,x])$ and $mmc^{-1}([x,\infty))$ are closed in Φ.*

Proof. For ease of presentation, we adopt the following notation: for any $x \in \Re$, $\mathcal{C}_x = c^{-1}(x)$. Take $x \in (0, \sup\iota(\Upsilon))$; to show that $mmc^{-1}([0,x])$ is closed, suppose not, and suppose that $A \in \left(\overline{mmc^{-1}([0,x])} \backslash mmc^{-1}([0,x])\right) \cap \Phi$. Let $A_n \in mmc^{-1}([0,x])$ be a sequence with $A_n \to A$. It follows from representation (2) and the definition of mmc that, for every $n \in \mathbb{N}$, every $f_n \in \sup(A_n, u, \{p_\Xi\})$ and every $x' > x$, there exists $p, q \in \mathcal{C}_{x'}$ and $g, h \in A_n$ such that $\sum_{s\in S} u(f_n(s)).p(s) < \sum_{s\in S} u(h(s)).p(s)$ and $\sum_{s\in S} u(f_n(s)).q(s) > \sum_{s\in S} u(g(s)).q(s)$. It follows from the representation and the fact that $A \in \left(\overline{mmc^{-1}([0,x])} \backslash mmc^{-1}([0,x])\right) \cap \Phi$ that for some $x' > x$ and all $f \in \sup(A, u, \{p_\Xi\})$, $\sum_{s\in S} u(f(s)).p(s) \geqslant \sum_{s\in S} u(h(s)).p(s)$ for all $p \in \mathcal{C}_{x'}$ and all $h \in A$. Since $A_n \to A$, it follows from the continuity of the representation that the hyperplane $u \circ f - u \circ h = 0$ in $\Delta(S)$ supports $\mathcal{C}_{x'}$, for at least one $h \in A$ and $f \in \sup(A, u, \{p_\Xi\})$. If this holds for some $x' > x$, it must also hold for any $x' > x'' > x$ by the same reasoning; it follows from the fact that Ξ is a nested family that there exists a support point of $\mathcal{C}_{x'}$ that is also a support point of $\mathcal{C}_{x''}$. Moreover, since $A \in \Phi$, such a point is not a support point of $\overline{\bigcup_{\mathcal{C}'\in\Xi} \mathcal{C}'}$. Since support points are those that do not belong to the relative interiors of the appropriate sets, it follows that $(C_{x''} \cap (ri(C_{x'}))^c) \cap ri(\overline{\bigcup_{\mathcal{C}'\in\Xi} \mathcal{C}'}) \neq \varnothing$, contradicting the strictness of Ξ. So $mmc^{-1}([0,x])$ is closed in Φ. The closeness of $mmc^{-1}([x,\infty])$ in Φ is a straightforward consequence of the continuity of the unanimity representation.

□

C.2 Proofs of results in Appendix A

Proof of Proposition A.1. The 'if' direction is straightforward. The 'only if' direction is a simple corollary of the proof of Theorem 1. On the one hand, if γ^1 and γ^2 are confidence equivalent, they yield identical choices over menus consisting entirely of constant acts; hence the utilities are the same up to positive affine transformation. On the other hand, if they are confidence equivalent, the family of functions $\{\gamma_r | r \in \mathcal{S}^f\}$ defined in the proof of Theorem 1 are the same, and so the confidence rankings are the same. □

Proof of Proposition A.2. Let the assumptions of the Proposition be satisfied and let (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) represent γ^1 and γ^2 respectively. Beyond using some terminology introduced in Section 3.2.3, the following notation shall prove useful. For any $A \in \wp(\mathcal{A})$, define $\mathcal{C}_A \in \Xi$ as follows: $\mathcal{C}_A = \max\{\mathcal{C} \in \Xi \mid \sup(A, u, \mathcal{C}) \neq \varnothing\}$. We consider the three parts in turn.

Part (i). The right-to-left implication is straightforward, so we only consider the left-to-right direction. It follows from representation (2) that $c_1(D_1(A)) = u(\iota_1(A))$ for all $A \in \Upsilon_1$ and similarly for decision maker 2. Since γ^1 is more decision averse than γ^2, $\iota_2(A) \leqslant \iota_1(A)$ for all $A \in \wp(\mathcal{A})$, and hence in particular $A \in \Upsilon_2$ implies $A \in \Upsilon_1$. It follows that, for each $A \in \Upsilon_2$, $c_1(D_1(A)) \geqslant c_2(D_2(A))$. This property extends to $\overline{\Upsilon_2}$ by the continuity of the D and c.

Define $D_3 : \wp(\mathcal{A}) \rightarrow \Xi$ by $D_3(A) = D_2(A)$ whenever $A \in \overline{\Upsilon_2}$ and $D_3(A) = \min\{D_1(A), D_2(A)\}$ otherwise. By definition, and the fact that for $A \in \overline{\Upsilon_2}$, $c_1(D_1(A)) \geqslant c_2(D_2(A))$, $c_2 \circ D_3$ satisfies the ordering conditions with respect to $c_1 \circ D_1$ in the Proposition. We now show that it is a cautiousness coefficient and that (u, Ξ, D_3, c_2) represents γ^2. Continuity of D_3 off the boundary of Υ_2 follows from the continuity of D_1, D_2 and the minimum. To establish continuity for menus on the boundary, it suffices to show that for every $A \in \overline{\Upsilon_2} \backslash \Upsilon_2$, $D_2(A) \subseteq D_1(A)$. For every $A \in \overline{\Upsilon_2} \backslash \Upsilon_2$, it follows from representation (2) that $D_2(A) = \mathcal{C}_A$; since, as noted above, $\Upsilon_2 \subseteq \Upsilon_1$, then either $A \in \Upsilon_1$, in which case $D_1(A) \supset D_2(A)$ by the definition of Υ, or $A \in \overline{\Upsilon_1} \backslash \Upsilon_1$, in which case $D_1(A) = \mathcal{C}_A = D_2(A)$. So $D_2(A) \subseteq D_1(A)$, and D_3 satisfies continuity. Extensionality and Φ-richness of D_3 follow from the extensionality and Φ-richness of D_1 and D_2, so D_3 is a cautiousness coefficient. Finally, we show that (u, Ξ, D_3, c_2) represents γ_2. Since, by definition, $D_3(A) = D_2(A)$ for all $A \in \overline{\Upsilon_2}$, $c_2 \circ D_3 = c_2 \circ D_2$ on $\overline{\Upsilon_2}$. Moreover, for all $A \notin \overline{\Upsilon_2}$, $\sup(A, u, D_3(A)) \neq \varnothing$, and so, by representation (2) (and the strictness of the confidence ranking), $\sup(A, u, D_3(A)) = \sup(A, u, D_2(A))$. Since (u, Ξ, D_2, c_2) represents γ_2, so does (u, Ξ, D_3, c_2), as required.

Part (ii). By the representation (2), for any $A \in \Phi_1$, $mmc_1(A) = c_1(\mathcal{C}_A)$ and similarly for decision maker 2. Moreover that, by a simple supporting hyperplane argument, for every $\mathcal{C} \in \Xi\backslash\{\overline{\bigcup_{\mathcal{C}'\in\Xi}\mathcal{C}'}\}$, there exists $A' \in \Phi_1$ such that $\mathcal{C}_{A'} = \mathcal{C}$. (It suffices to consider any supporting hyperplane of $\mathcal{C}$, and take the menu containing an act and a constant act corresponding to the hyperplane, as in Lemmas C.8–C.10, for example.) Note finally that, since $A \in \Phi_1$ iff $\sup(A, u, \overline{\bigcup_{\mathcal{C}'\in\Xi}\mathcal{C}'}) \neq \emptyset$, it follows from the fact that γ^1 and γ^2 are represented by the same u and Ξ that $\Phi_1 = \Phi_2$. It follows from representation (2) that γ^1 is less cost motivated than γ^2 iff $mmc_1(A) \geqslant mmc_2(A)$ for all $A \in \Phi_1$. By the first remark above, this holds iff $c_1(\mathcal{C}_A) \geqslant c_2(\mathcal{C}_A)$ for all $A \in \Phi_1$; by the second remark, this holds iff $c_1(\mathcal{C}) \geqslant c_2(\mathcal{C})$ for all $\mathcal{C} \in \Xi\backslash\{\overline{\bigcup_{\mathcal{C}'\in\Xi}\mathcal{C}'}\}$. By the continuity of c_1 and c_2, this holds iff $c_1(\mathcal{C}) \geqslant c_2(\mathcal{C})$ for all $\mathcal{C} \in \Xi$, yielding the required equivalence.

Part (iii). By the representation (2), for any $A \in \Upsilon_1$, $\iota_1(A) = c_1(D_1(A))$ and similarly for decision maker 2. By representation (2), γ^1 is more motivation-calibrated decision averse than γ^2 iff for all $A, B \in \wp(\mathcal{A})$, $\iota_1(A) \leqslant mmc_1(B) \Rightarrow \iota_2(A) \leqslant mmc_2(B)$. Given the observations made in the proof of part (ii), this holds iff, for all $A \in \Upsilon_2$ and $\mathcal{C} \in \Xi\backslash\{\overline{\bigcup_{\mathcal{C}'\in\Xi}\mathcal{C}'}\}$, $c_1(D_1(A)) \leqslant c_1(\mathcal{C}) \Rightarrow c_2(D_2(A)) \leqslant c_2(\mathcal{C})$. Since c_1 and c_2 are order-preserving and -reflecting, this holds iff $D_1(A) \subseteq \mathcal{C} \Rightarrow D_2(A) \subseteq \mathcal{C}$ for all $A \in \Upsilon_2$ and $\mathcal{C} \in \Xi\backslash\{\overline{\bigcup_{\mathcal{C}'\in\Xi}\mathcal{C}'}\}$, and hence iff $D_1(A) \supseteq D_2(A)$ for all $A \in \Upsilon_2$. So γ^1 is more motivation-calibrated decision averse than γ^2 iff $D_1(A) \supseteq D_2(A)$ for all $A \in \Upsilon_2$. Note moreover that if γ^1 is more motivation-calibrated decision averse than γ^2, then $\Upsilon_2 \subseteq \Upsilon_1$: for any $A \in \Upsilon_2$, $D_2(A) \supset \mathcal{C}_A$, and hence, since γ^1 is more motivation-calibrated decision averse, $D_1(A) \supset \mathcal{C}_A$, so $A \in \Upsilon_1$.

It follows immediately from representation (2) and the uniqueness conditions in Theorem 1 that the existence of a cautiousness coefficient satisfying the conditions in part (iii) implies that γ^1 is more motivation-calibrated decision averse than γ^2. To show the left-to-right direction of part (iii), we show that if $\Upsilon_2 \subseteq \Upsilon_1$ and $D_1(A) \supseteq D_2(A)$ for all $A \in \Upsilon_2$, then there exists cautiousness coefficient D_3 representing γ^2 such that $D_3(A) \subseteq D_1(A)$ for all $A \in \wp(\mathcal{A})$. To this end, let $D_3(A) = \min\{D_1(A), D_2(A)\}$ for all $\mathcal{A}$. This is a cautiousness coefficient: extensionality, continuity and Φ-richness follow from the extensionality, continuity and Φ-richness of D_1 and D_2, the continuity of the minimum, and the fact that

$D_1(A) \supseteq D_2(A)$ on Υ_2. By definition, $D_3(A) \subseteq D_1(A)$ for all $A \in \wp(\mathcal{A})$. By a similar argument to that used in the proof of part (i), (u, Ξ, D_3, c_2) represents γ^2. So (u, Ξ, D_1, c_1) and (u, Ξ, D_3, c_2) satisfy the required conditions.

□

Proposition C.3. *Let γ^1 and γ^2 satisfy axioms A1–A10 and be confidence equivalent.*

- *If does not follow from γ^1 being more decision averse and less cost motivated than γ^2 that γ^1 is more motivation-calibrated decision averse than γ^2.*
- *It does not follow from γ^1 being more decision averse and more motivation-calibrated decision averse than γ^2 that γ^1 is less cost motivated than γ^2.*

Proof. Consider any (u, Ξ, D_2, c_2) satisfying the conditions in Theorem 1 and let γ^2 be the choice correspondence it represents. Take two continuous functions $f_D : c_2(\Xi) \to c_2(\Xi)$ and $f_c : c_2(\Xi) \to \Re$ such that f_D is surjective and f_c is strictly increasing. Define $D_1 : \wp(\mathcal{A}) \to \Xi$ by $D_1(A) = c_2^{-1}(f_D(c_2(D_2(A))))$ and $c_1 : \Xi \to \Re_{\geqslant 0}$ by $c_1(\mathcal{C}) = f_c(c_2(\mathcal{C}))$. It is straightforward to check that these are well-defined cautiousness coefficients and cost functions respectively; let γ^1 be the choice correspondence represented by (u, Ξ, D_1, c_1). Let us say that a function $\Gamma : X \to Y$, where $X \subseteq Y$, is *upper-valued* if, for every $x \in X$, $f(x) \geqslant x$. Note that, by Proposition A.2 (and the order-reflecting and preserving properties of the cost function), γ^1 and γ^2 are ordered by decision aversion iff $f_c \circ f_D$ is upper-valued, they are ordered by motivation-calibrated decision aversion iff f_D is upper-valued, and they are ordered by cost motivation iff f_c is upper-valued.

To establish the first part of the proposition, it thus suffices to find f_c and f_D such that f_c is upper-valued, f_D is not, but $f_c \circ f_D$ is. Let $a = \inf c_2(\Xi)$ and $b = \sup c_2(\Xi)$, and consider $f_D(x) = \frac{1}{b-a}(x-a)^2 + a$, and $f_c(x) = x + \frac{b-a}{2}$. It is straightforward to check that f_D is not upper-valued, whilst f_c and $f_c \circ f_D$ are, yielding the required example. Similarly, taking for example $f_D(x) = (b-a)^{\frac{1}{2}}(x-a)^{\frac{1}{2}} + a$, and $f_c(x) = \frac{3}{4}(x-a) + a$ for $x \leqslant \frac{b-a}{4} + a$, and $f_c(x) = \frac{13}{12}(x-a) - \frac{b-a}{12} + a$ for $x \geqslant \frac{b-a}{4} + a$ yields an example establishing the second part of the proposition.

□

References

AKERLOF, G. A. (May, 1991): "Procrastination and Obedience," *The American Economic Review*, 81(2), 1–19.

ALIPRANTIS, C. D., AND K. C. BORDER (2007): *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edn.

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

BEWLEY, T. F. (1986 / 2002): "Knightian decision theory. Part I," *Decisions in Economics and Finance*, 25, 79–110.

BILLINGSLEY, P. (2009): *Convergence of Probability Measures*. John Wiley & Sons.

BUTURAK, G., AND Ö. EVREN (2014): "A Theory of Choice When "No Choice" is an Option," Discussion paper, NES.

DANAN, E. (2003a): "A behavioral model of individual welfare," Discussion paper, Université Paris 1.

——— (2003b): "Revealed Cognitive Preference Theory," Discussion paper, EUREQua, Université de Paris 1.

DEKEL, E., B. L. LIPMAN, AND A. RUSTICHINI (2001): "Representing Preferences with a Unique Subjective State Space," *Econometrica*, 69, 891–934.

DHAR, R. (1997): "Consumer Preference for a No-Choice Option," *The Journal of Consumer Research*, 24(2), 215–231.

ERGIN, H., AND T. SARVER (2010): "A unique costly contemplation representation," *Econometrica*, 78(4), 1285–1339.

GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): "Differentiating ambiguity and ambiguity attitude," *J. Econ. Theory*, 118(2), 133–173.

GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin expected utility with non-unique prior," *J. Math. Econ.*, 18(2), 141–153.

GRUBER, M. J. (1996): "Another puzzle: The growth in actively managed mutual funds," *The journal of finance*, 51(3), 783–810.

GUL, F., AND W. PESENDORFER (2001): "Temptation and Self-Control," *Econometrica*, 69(6), 1403–1435.

HART, K. P., J.-I. NAGATA, AND J. E. VAUGHAN (2004): *Encyclopedia of general topology*. Elsevier.

HILL, B. (2012): "Confidence in preferences," *Social Choice and Welfare*, 39(2), 273–302.

——— (2013): "Confidence and decision," *Games and Economic Behavior*, 82, 675–692.

——— (2014): "Incomplete preferences and Confidence," Discussion paper, HEC Paris.

HOLMSTROM, B. (1984): "On the Theory of Delegation," in *Bayesian Models in Economic Theory*, ed. by M. Boyer, and R. Kihlstrom. North-Holland, New York.

KELLEY, J. L. (1975): *General Topology*. Springer, New York.

KOOPMANS, C. (1964): "On the flexibility of future preferences," in *Human judgements and rationality*, ed. by M. Shelley, and G. Bryan. John Wiley and Sons.

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

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

MALKIEL, B. G. (1995): "Returns from investing in equity mutual funds 1971 to 1991," *The Journal of Finance*, 50(2), 549–572.

MARSCHAK, J., AND K. MIYASAWA (1968): "Economic Comparability of Information Systems," *International Economic Review*, 9(2), 137–174.

MASATLIOGLU, Y., AND E. A. OK (2005): "Rational choice with status quo bias," *Journal of Economic Theory*, 121(1), 1–29.

O'DONOGHUE, T., AND M. RABIN (2001): "Choice and Procrastination," *Quart. J. Econ.*, 116(1), 121–160.

ORTOLEVA, P. (2010): "Status quo bias, multiple priors and uncertainty aversion," *Games and Economic Behavior*, 69(2), 411–424.

——— (2013): "The price of flexibility: Towards a theory of Thinking Aversion," *Journal of Economic Theory*, 148(3), 903–934.

RIELLA, G., AND R. TEPER (2014): "Probabilistic dominance and status quo bias," *Games and Economic Behavior*, 87, 288–304.

SAVOCHKIN, A. (2014): "Mistake Aversion and a Theory of Robust Decision Making," Discussion paper, Collegio Carlo Alberto.

SEN, A. K. (1971): "Choice Functions and Revealed Preference," *Rev. Econ. Stud.*, 38(3), 307–317.

STIGLER, G. J. (1961): "The Economics of Information," *The Journal of Political Economy*, 69(3), 213–225.

TVERSKY, A., AND E. SHAFIR (1992): "Choice Under Conflict: The Dynamics of Deferred Decision.," *Psychological Science (Wiley-Blackwell)*, 3(6), 358–361.