A Representation Result for Choice under Conscious Unawareness *

Oliver Walker†
Grantham Research Institute on Climate Change and the Environment
London School of Economics and Political Science

Simon Dietz
Grantham Research Institute on Climate Change and the Environment
and Department of Geography and Environment
London School of Economics and Political Science

November 22, 2012

Abstract

There are many examples in policy-making, investment and day-to-day life where the set of contingencies the decision-maker is aware of does not resolve all uncertainty about the consequences of actions. In such circumstances, the decision-maker may nevertheless reason that there exist certain aspects of the 'full' state space of which she is unaware, that is, she may think it is possible she is unaware of something. We call this type of belief conscious unawareness and claim that its presence may lead to a violation of Savage's Sure Thing Principle for reasons parallel to those at play in the Ellsberg examples, but differently motivated. We then specify a choice setting in which the primary domain of choice is a set of actions stated naturally in English, but where the decision-maker also has preferences on the set of derivative actions. A derivative action maps from the set of permutations – the product space of the set of contingencies she can conceive of (her subjective state space) and the set of payoff assignments to the actions – to a space of consequences. We obtain a representation result that makes choice, in cases where conscious unawareness is a major concern, tractable by means of some of the standard analytical tools of risk and ambiguity analysis (in particular, those developed in light of Klibanoff, Marinacci, and Mukerji's (2005) model). The representation allows us

*Our research has been supported by the Grantham Foundation for the Protection of the Environment, as well as the Centre for Climate Change Economics and Policy, which is funded by the UK's Economic and Social Research Council (ESRC) and by Munich Re. The usual disclaimers apply.

†Email address for correspondence: o.j.walker@lse.ac.uk

to characterise the decision-maker's attitude towards perceived payoff uncertainty arising from factors she is unaware of; in particular, we characterise ignorance aversion. Using the same framework, we are also able to state a more general representation that allows us to capture source preference in examples where the decision-maker is consciously unaware.

Keywords: Unawareness, Ignorance, Conscious Unawareness, Ambiguity, Uncertainty, Ignorance Aversion, Ambiguity Aversion, Non Expected Utility, Source Preference, "Small Worlds", Climate Change Policy

JEL Classification Numbers: D81

1 Introduction

?'s (**?**) theory of subjective expected utility (SEU) posits a decision-maker (DM) with ready access to a *full state space*, comprising all possible "descriptions of reality, leaving no relevant feature undescribed". The contingencies that make up the full state space are so finely described that, whatever action the DM takes, any uncertainty about what consequence that action might lead to is resolved by the state that transpires and, furthermore, the DM knows this to be the case.

There are, however, many examples in policy-making, investment and day-to-day life where the contingencies the DM takes into account do not resolve all payoff uncertainty in this manner. Consider, for example, a trader speculating on the price of oil. Given the large number of factors that determine the price of oil and the complicated manner in which they interact, it seems highly doubtful that the trader would be able to describe all of the states relevant to her payoff as required by Savage's framework. Rather, it is likely that the set of contingencies the trader does take into account – what we call her *subjective state space* – omits certain relevant details or distinctions and thus does not resolve all of the payoff uncertainty faced. Where this is so, we say the trader is *unaware* of the full state space.

It is widely accepted that unawareness is a pervasive feature of decision problems and for this reason, it is also commonplace for DMs to reason about the possibility that there exist relevant considerations of which they are unaware. We describe a DM who believes she may be unaware of the full state space – and hence who regards her own subjective state space as offering a possibly incomplete account of the payoff-relevant uncertainty she faces – as *consciously unaware.*

Savage's theory was evidently not designed for consciously unaware agents and in our view such agents may be justified in violating SEU. In particular, we believe that allowing for conscious unawareness calls into question the Sure Thing Principle (STP), a necessary condition for Savage's representation. The following example illustrates our point.

Example 1: Suppose the oil trader has the opportunity to perform the action a, given as "Spend \$1 million on six-month oil futures contracts at \$100 per barrel". She knows that a's payoff will be determined by the oil price in six months' time, p, but can only imagine two contingencies that might have a bearing on this. These are s_1, "peace holds in the Middle East", or s_2, "war breaks out in the Middle East". If s_1 occurs, she thinks p could take any value between \$70 and \$115, while if s_2 occurs, she thinks it could be anything between \$85 and \$130.

She may also take herself to the races, where she will be able to gamble on the performance of the horse "Mighty Monty" in a race. The payoffs from these gambles depend on whether or not Mighty Monty wins and, once again, she can think of two contingencies that could influence this. These are s_3, "Mighty Monty has fully recovered from a bout of tendinitis" – in which case she believes he is bound to win his race – and s_4, "Mighty Monty has not fully recovered from tendinitis" – in which case she thinks he is sure to lose. Suppose the only contingencies she ever takes into account when evaluating different actions is given by the product space, $S := \{s_1, s_2\} \times \{s_3, s_4\}$.

The trader recognises that S does not resolve all payoff uncertainty she faces – in particular, any state in S seems consistent with a returning a wide range of payoffs. She is therefore consciously unaware, but this does not affect all of her choice set in the same way. For, in her view, while there is no state in S that resolves the payoff uncertainty pertaining to a, every state resolves the payoff uncertainty pertaining to gambles on Mighty Monty. We will say that she thus *understands* gambling on Mighty Monty, but not purchasing oil futures.

Our theory allows this asymmetry to be reflected in preferences in a manner that is inconsistent with SEU. To see this, suppose the trader is presented with a choice between the following "derivative" action:

$$c_1 = \begin{array}{c}\text{If } a \text{ pays out more than \$1.1 million} \\ \text{receive \$100, otherwise receive \$0}\end{array}$$

and the gamble c' given as follows:

$$c_2 = \begin{array}{c}\text{If Mighty Monty wins the race} \\ \text{receive \$100, otherwise receive \$0}\end{array}$$

and imagine the trader reports a strict preference for c_2 over c_1. Where $x - b$ refers to a prospect that yields \$100 less whatever the prospect b pays out, suppose the trader is then asked to choose between $100 - c_1$ to $100 - c_2$. The payoff structures of the four prospects a', a^*, $100 - c$, and $100 - c'$ are illustrated in the table below where MM refers to the contingency "Mighty Monty wins".

	Dollar payoff in case			
	MM	MM	$\neg MM$	$\neg MM$
	$p > 110$	$p \leq 110$	$p > 110$	$p \leq 110$
c_1	100	0	100	0
c_2	100	100	0	0
$100 - c_1$	0	100	0	100
$100 - c_2$	0	0	100	100

Given her preference for c_2 ahead of c_1, Savage's STP requires the trader to prefer $100-c_1$ to $100-c_2$. This is typically justified along the following lines: c_1 and c_2 offer identical payoffs in the first and fourth columns of Table 1, so strict preference for c_2 over c_1 should imply strict *conditional preference* for c_2 over c_1 *given* "MM and $p \leq 110$" or "$\neg MM$ and $p > 110$". But this means that $100 - c_1$ must be conditionally preferred to $100 - c_2$ given the same information (the two prospects' respective conditional payoff structures are identical to those of c_2 and c_1). And since $100 - c_1$ and $100 - c_2$ offer the same payoffs under any other contingency, this conditional preference should dictate the trader's unconditional preference for $100 - c_1$ over $100 - c_2$.

However, we wish allow the trader to strictly prefer c_2 to c_1 *and* $100-c_2$ to $100 - c_1$, in violation of the STP and hence of SEU. The trader may reason that she has a general preference for taking on actions she understands, and thus that a relevant feature of the two decision problems – overlooked by the argument above – is that she understands c_2 and $100-c_2$, but not c_1 and $100-c_1$. She might therefore reject the claim that preferring c_2 to c_1 commits her to a conditional preference for c_2 over c_1 given "MM and $p \leq 110$" or "$\neg MM$ and $p > 110$", arguing instead that it reflects her aversion to actions she does not understand. This gives her grounds to reject the the STP in this case and choose $100 - c_2$ ahead of $100 - c_1$.

We propose an alternative to SEU that is consistent with cases such as this. In our framework, the DM is endowed with a subjective state space, S, and knows that any action she might carry out will lead to a consequence within a given space X. The choice set of primary interest is then a set of actions, $\mathcal{A}$, given as sentences in English describing things to do such as "Spend \$1 million on six-month oil futures at \$100 per barrel".

To reveal how the DM conceives of the members of $\mathcal{A}$, we suppose she has preferences over prospects akin to the "derivative" actions described in Example 1. These are defined by introducing, for each $s \in S$, the set W_s

consisting of all maps from $\mathcal{A}$ to X. W_s is interpreted as the list of every possible profile of payoffs $\mathcal{A}$ might induce if the subjective state s were to occur: it is the set of *permutations* under s. If the DM did not think state s resolved all of the payoff uncertainty pertaining to action a – what we term not understanding action a – then she would be willing to gamble on multiple permutations in W_s that assigned different payoffs to a. We interpret the DM's willingness to gamble on a permutation in W_s as the same as her regarding the payoff profile it stands for as possible if s occurs.

Our representation is obtained by applying familiar regularity conditions over various choice sets. First, we assume that her preferences over the set of Anscombe-Aumann acts defined on S are consistent with SEU and thus encode a unique subjective prior on S, π, and a utility function, v, that represents her attitude to risky gambles on X. This implies that the DM's choices are consistent with SEU over the set of actions she understands. Second, we assume the DM has conditional preferences on maps from W_s to X for every s. Imposing well-known assumptions, we can obtain a subjective prior, μ_s, on each W_s and a utility function ϕ on X, such that we find action a is preferred to a' if and only if:

$$\sum_{s\in S}\pi(s)u\left(\int_{W_s}\phi(w(a))\mathrm{d}\mu_s\right) \;\geq\; \sum_{s\in S}\pi(s)u\left(\int_{W_s}\phi(w(a'))\mathrm{d}\mu_s\right) \qquad (1)$$

where $u := v\circ\phi^{-1}$ and we write $w(a)$ for the consequence action a produces in the payoff profile w.

The form of the representation in (1) and the way we derive it from our assumptions are familiar from **?**'s (KMM, **?**) smooth model of decision under ambiguity, but the motivation and structural setting underpinning our result is quite different. Our DM does not depart from SEU because she faces ambiguity over the true probability density function (pdf) over the state space; rather, our DM believes that the subjective state space she has in mind is insufficiently rich to identify every action's payoff and hence that there are actions in her choice set that she does not understand. When choosing from a choice set that includes some actions she understands and some that she does not, she may wish to exercise particular caution (or recklessness) over the actions she does not understand, and hence violate SEU. We discuss the connection between this sort of behaviour and ambiguity aversion below.

We hope that our representation will be particularly helpful in various policy settings where the fact that there is unawareness is a major concern. One such domain is policy on climate change. Here, the state of scientific knowledge about the links between emissions of greenhouse gases and changes to physical climate variables such as temperature, precipitation and sea level is recognised to be far from exhaustive (**?**; **?**), and our understanding of the interface between the climate and the economy is thought to be similarly incomplete (e.g. **?**; **?**; **?**). Under such circumstances, some of the states we envisage – even described at the most minute level of detail we can conceive of – seem consistent with almost any payoff, no matter what climate policy we pursue. This means not only that (in our view) conscious unawareness should be a significant consideration in climate policy, but also that the problem is very difficult to analyse using existing decision models (including those that can accommodate conscious unawareness). Our theory makes choice problems such as these analytically tractable.

To illustrate, suppose that a policy-maker's subjective state space consists of the contingencies s and s', where:

$$s = \text{"Global temperature depends sensitively on the atmospheric concentration of greenhouse gases"}$$
$$s' = \text{"Global temperature does not depend sensitively on the atmospheric concentration of greenhouse gases"}$$

and that she has a choice between the following actions:

$$a = \text{"Cut greenhouse gas emissions by 50\% by 2050"}$$
$$a' = \text{"Cut greenhouse gas emissions by 30\% by 2050"}$$

Consider the task of conducting an economic evaluation of these climate-policy actions (i.e. a cost-benefit analysis), in which the set of consequences is just a range of possible monetary outcomes. Given our degree of understanding of the problem, it seems reasonable to allow that both a and a' could pay out any amount in X under both s and s'. Other representations of choice under conscious unawareness (for example **?**, **?**) do not allow the DM to hold beliefs about the relative likelihood of either action paying any given consequence under either of the states. They therefore require the policy maker to regard a and a' as equally good in both states and, given any action, that she is indifferent between which of the two states does transpire. Yet it seems obvious that the policy maker would regard state s as "bad news" under either action and that, given s, the policy maker would prefer to carry out a over a'. Such preferences are consistent with our theory and would imply that $\mathbb{E}_{\mu_s}[\phi(w(a))] > \mathbb{E}_{\mu_s}[\phi(w(a'))]$ and $\mathbb{E}_{\mu_{s'}}[\phi(w(a))] > \mathbb{E}_{\mu_s}[\phi(w(a))]$.

Another advantage of allowing beliefs over W_s and $W_{s'}$ is that we can capture DMs' aversion (or predilection) towards less well-understood actions in a familiar fashion. Consider two ways to reduce the atmospheric concentration of greenhouse gases. The first, b, involves the replacement of fossil-fuel power plants with renewables, such as onshore wind farms. The second, b', involves the use of a 'geo-engineering' technique, whereby iron is poured into the oceans, in order to stimulate blooms of phytoplankton, which remove carbon dioxide from the atmosphere. Suppose that under the policy-maker's μ_s, the expected net monetary benefits of b are equal to b', but that the pdf on X entailed by b' and μ_s is a mean-preserving spread of that entailed by b and μ_s. One might say that the policy-maker "better understands" b than b' given the occurrence of state s. Always preferring actions over less well understood alternatives with the same expected payoff – what we call *ignorance aversion* – is equivalent to the concavity of the function ϕ in our framework. This mirrors exactly the characterization of risk aversion in SEU theory and ambiguity aversion in KMM's approach.

Source Preference here!

The rest of this paper is organised as follows. First, we introduce the elements of the choice setting and the DM's preferences, before setting out our assumptions and result. Then we give behavioural characterisations of "ignorance aversion" and "more ignorance averse", showing that these are formally equivalent to concavity properties of the function ϕ. In Section 4 we set out a somewhat generalised version of our representation that can accommodate "source preference", before ending with a discussion of our assumptions and the connection between this work and that on ambiguity. All proofs are in the Appendix.

2 Subjective State Spaces and Choice

In Savage's classic framework, a DM's preferences are given over a set of maps from states of the world to consequences called *Savage acts*. Every element of Savage's state space encodes all relevant information about the state of the world in which consequences are received, while each consequence gives a full account of all the things that matter to the DM. Conceptualising the choice set is therefore a trivial exercise for the DM: for any choice, the set of contingencies her payoff depends on and the payoff received under each contingency is given explicitly.

We wish to allow for cases, such as that of Example 1, where the DM may have difficulty understanding the choice set in the way Savage's framework presupposes. We thus model the DM as choosing from a set of *actions* composed of simple descriptions of things to do in English such as "Spend \$1 million on six-month oil futures at \$100 per barrel" or "Purchase a £10 bet on Mighty Monty". Clearly, it is possible for a DM to choose from these kinds of prospects without knowing what the payoff-relevant contingencies are.

How might a DM faced with such a choice set formulate her decision? We offer one account here to help the reader interpret the formal setting and restrictions of the next section. This is intended only to be suggestive and we emphasise that alternative accounts may be applicable.

Suppose the DM knows that any consequence she might receive has some *ex* *ante* monetary equivalent on a bounded interval, X, and she knows that, for every action, there exists a "correct" description of it as a map from states of the world to X (that is, as a Savage act). This correct characterisation could be the way a hypothetical omniscient analyst would think of the action. Furthermore, the DM knows that, if only she could describe the set of payoff-relevant contingencies in sufficient detail, she would understand all of the actions in the choice set in this correct fashion. We can think of this level of detail as corresponding to a set of propositions, P, that are assigned truth or falsehood in every state and refer to the set of states described in this way as *the full state space*, Ω.

Our decision maker only considers things at a certain level of detail – in other words, she takes a limited set of propositions, Q, into account when formulating her decision. She knows this to be the case and might further suspect that the level of detail is less than that of the full state space (that is, she might suppose that $Q \subset P$ without having any idea of the composition of $P \backslash Q$). But she nonetheless thinks there are some actions she does *understand* in the following sense. She believes there is a collection of propositions that she takes into account, $Q' \subseteq Q$, and a set of contingencies described at the level of detail of Q', $S_{Q'}$, such that, under their respective truth assignments to Q' and P, each member of $S_{Q'}$ implies some subset of Ω and each member of Ω implies precisely one member of $S_{Q'}$. She then believes she understands a if, whenever a – in its "correct" rendering as a Savage act – pays out $x \in X$ in state $\omega \in \Omega$, she is certain that a pays out x in $s_{Q'} \in S_{Q'}$ where $s_{Q'}$ is implied by ω. Where this is so we say that the DM understands a under Q'.

Let A^* refer to the set of actions the DM understands under any Q' and, for every $a \in A^*$, write Q_a for the $\supseteq$-maximal set of propositions under which the DM understands a. Define Q^* as $\bigcup_{a \in A^*} Q_a$ and make the assumption that the DM understands every $a \in A^*$ under Q^*. We call S_{Q^*} – henceforth denoted S – the DM's *subjective state space*[1] and say Q^* is the set of propositions she is *aware of*. If the DM does not believe she understands every action – that is, if A^* does not encompass the full set of actions – then we describe the DM as being *consciously unaware*.

It follows from this that if the DM is consciously unaware and a is an action she believes she does not understand, she cannot be certain what a pays out under all members of S. That is, there must be some $s \in S$ and an $x \in X$ such that she'd pay money for a token entitling her to \$100 in case, conditional on s, a yielded at least x *and* she'd pay money for a token that paid \$100 if, conditional on s, a yielded less than x. Our framework reveals the DM's conscious unawareness in precisely this way by considering her

[1]We make a number of non-trivial technical steps here: why is A^* non-empty? why should the $\supseteq$-maximal set of propositions under which the DM understands a exist? why should the DM understand every $a \in A^*$ under Q^*? For more detail on these points see ?.

preferences over a choice set made up of bets on the payoff performance of the actions conditional on subjective states. These bets are what we call *derivative actions*.

One possible objection to this story concerns its applicability to our motivating example. In Example 1 the oil trader knows that the payoff of action c_1 is determined by the realisation of p: it might thus be countered that the DM in this case *can* understand c_1 as a Savage act – namely one that maps from the value of p to X – and that she is therefore not consciously unaware in the sense set out in this section. The DM, however, may reason as follows: although it is clear what payoff c_1 would deliver conditional on any value of p, it is not clear that every p is possible in a *logically consistent* state of affairs; therefore it *need not be the case* that a simple map from p to X is consistent with c_1's correct characterisation as a Savage act (which uses only logically consistent states of the world). Put another way, the DM regards the statement "$p = 110$" as just another way of saying "Action c pays out \$1.1 million", rather than a description of a member of the full state space. On this basis the DM may harbour reasonable doubts over whether she does understand c_1 and thus be consciously unaware in the example.

A second objection is more general. We state above that the DM believes that there is a set of propositions, P, with the property that awareness of them implies the ability to understand the full action set correctly, but one may question whether mere awareness of all relevant considerations is sufficient for this. Could it not be the case, for example, that the DM was aware of every member of P but was unable to correctly process the logical connections that hold between the members of P and the sets of actions and payoffs, and hence unable to understand all of the actions correctly? There are two ways of countering this objection. First, we may insist that part of what it is to be aware of a set of propositions (or actions and payoffs) is to perceive correctly all logical connections between them: thus awareness of P would commit the DM to knowledge of what each action pays out in every state, as this information is implied by all consistent truth assignments to P. This is the approach adopted in much of the literature on modelling unawareness (for example, **?**, **?**, **?**), where it corresponds to the assumption of *awareness generated by primitive propositions*. An alternative response is to follow **?** in allowing the DM's ability to establish these connections to be dependent on her degree of awareness; in this case our account embeds the assumption that full awareness corresponds to a full ability to perceive logical connections between propositions, actions, and payoffs.

INSERT CITATION TO MUKERJI SOMEWHERE

3 The Representation

The DM chooses from a set of actions, denoted $\mathcal{A}$ with typical members written a, a'. Her preferences over $\mathcal{A}$ are represented by the binary relation $\succeq^*$ on $\mathcal{A}$ with respective asymmetric and symmetric components $\succ^*$ and $\sim^*$. $\succeq^*$ is the preference relation of primary interest to us, but we arrive at our representation indirectly by placing restrictions on the DM's preferences over different sets to $\mathcal{A}$ and then requiring $\succeq^*$ to be consistent with these other relations in a particular way.

To introduce these additional choice sets, first let there be a consequence space X with generic elements x, x', equal to some bounded interval on the real line, and use $\mathcal{B}$ to denote the Borel σ-algebra of X. One way to interpret X is as encompassing all the *ex ante* monetary valuations the DM might attribute to the result of an action as in Section 2 above. Such an interpretation would be consistent with investment choices, for example. Second, we assume our DM is endowed with a finite space, S – called her subjective state space with typical members s, s' – that is composed of every contingency she can conceive of. Write $\mathcal{E}$ for 2^S, the subjective event space.

Now define $\mathcal{G} := X^{\mathcal{A}}$ with typical elements g, g'. Endowing $\mathcal{G}$ with the topology of pointwise convergence, let W be the product space $\mathcal{G} \times S$ containing typical members w, w'. W is interpreted as the set of *permutations* that might arise: that is, it comprises every possible state combined with every possible payoff profile over $\mathcal{A}$. A permutation resolves all uncertainty over what member of S obtains and all payoff uncertainty pertaining to $\mathcal{A}$: so $w = \{s, g\}$ represents the permutation where s occurs and each action a pays out $g(a)$. For each $s \in S$, define $W_s : s \times \mathcal{G}$, the subspace of permutations in which s obtains, and let $\mathcal{B}_s$ be the Borel σ-algebra generated by the relative topology on W_s.

Using this, we can introduce a further choice set, dubbed the *space of derivative actions* and denoted $\mathcal{C}$. $\mathcal{C}$ is defined as the set of $\mathcal{B}$-measurable functions from W to X, with typical members of $\mathcal{C}$ called *derivative actions* and written c, c' etc. A derivative action is interpreted as a prospect that pays out some amount depending on what permutation obtains: that it, it is a bet whose payoff is determined by the realisation of the subjective state and the payoffs of all actions. If the subjective state s occurs and the payoff profile over $\mathcal{A}$ turns out to be g, then derivative action c pays out $c(w)$ where $w = \{s, g\}$. For clarity, we show how this formal structure could be used to describe the choice setting in Example 1.

Example 1, continued. We have the subjective state space S and an interval of possible monetary payoffs, X. For notational consistency, let a' be an action described just as c_2 is in the example and suppose the space of actions is just $\{a, a'\}$, $W_{\{s_1,s_3\}}$ is the set of permutations where

$\{s_1, s_3\}$ obtains. For instance, "$\{s_1, s_3\}$ is true, a pays out x, and a' pays out y" belongs to $W_{\{s_1,s_3\}}$ for any $x, y \in X$. Where $g(a) = x$ and $g(a') = y$, we use the compact notation $\{\{s_1, s_3\}, g\}$ for this permutation.

c_1 is thus a derivative action in this framework: under all permutations where a pays out more than \$1.1 million, it delivers \$1,000; in all others it pays nothing. Thus, where s can be any member of S, we have $c(w) = \$1,000$ for $w \in \{\{s, g\} : g(a) > \$1,100,000\}$ and $c(w) = \$0$ for all other w.

We suppose the DM is endowed with a preference relation $\succeq$ over $\mathcal{C}$ with respective asymmetric and symmetric components $\succ$ and $\sim$.

Our first restriction on $\succeq$ is a familiar independence condition. To state it, we use the notation $\{c, s; c'\}$ to refer to the derivative action c'' that satisfies $c''(w) = c(w)$ if $w = \{s, g\}$ for any $g \in \mathcal{G}$ and $c'' = c'(w)$ otherwise. In words, c'' pays out the same as c whenever s occurs and the same as c' under any other subjective state. The restriction is then as follows:

Axiom 1 (Monotonicity) *For any $c, c', c'' \in \mathcal{C}$, $c \succeq \{c', s; c\}$ iff $\{c, s; c''\} \succeq \{c', s; c''\}$.*

Given Monotonicity, we can define a conditional preference $\succeq_s$ for each $s \in S$ as $c \succeq_s c'$ iff $\{c, s; c''\} \succeq \{c', s; c''\}$ for any $c'' \in \mathcal{C}$. We say state s is *null* whenever $c \sim_s c'$ for all $c, c' \in \mathcal{C}$.

We call the next restriction on $\succeq$ an "Assumption" rather than an "Axiom" because its behavioural content is not immediate. More primitive behavioural conditions that are equivalent to it have been described by **?**.

Assumption 1 (Derivative-SEU) *Every $s \in S$ is either null or such that there exists a bounded, continuous, strictly increasing function $\phi_s : X \to \mathbb{R}$ and a probability measure on $\mathcal{B}_s$, denoted μ_s, such that:*

$$c \succeq_s c' \iff \int_{W_s} \phi_s(c(w))\mathrm{d}\mu_s \geq \int_{W_s} \phi_s(c'(w))\mathrm{d}\mu_s \tag{2}$$

for all $c, c' \in \mathcal{C}$, and there is at least one non-null s for which there exists an $E \in \mathcal{B}_s$ with $1 > \mu_s(E) > 0$.

The restriction that for one s there is some $E \in \mathcal{B}_s$ such that $1 > \mu_s(E) > 0$ implies that the DM does not regard the payoff of all actions as certain conditional on all non-null states. In the narrative of Section 2, this is the essence of conscious unawareness. It corresponds to **?**'s (**?**) characterisation of "the DM believes that if s occurs she may be unaware of something" as her willingness to gamble on multiple mutually inconsistent payoff profiles conditional on state s.

A striking feature of Derivative-SEU is that it requires $\succeq_s$ to satisfy Savage's STP for all s. Even in conjunction with our other restrictions, this does not imply that $\succeq$ satisfies the principle, something we argued against in Example 1. One may nonetheless wonder whether similar examples may be constructed that call Derivative-SEU into question. To allay such fears, we note that Example 1 "works" because the DM understands c_2 and $100 - c_2$ but not c_1 and $100 - c_1$; a parallel example would thus require the DM to understand $\{c, E; c''\}$ and $\{c', E; c'''\}$ but not $\{c, E; c'''\}$ and $\{c', E; c''\}$ for some $c, c', c'', c''' \in \mathcal{C}$ and $E \in \mathcal{B}_s$ and hence report $\{c, E; c''\} \succeq_s \{c', E; c''\}$ but $\{c', E; c'''\} \succ_s \{c, E, c'''\}$. However, it is easy to verify that this can only hold if there is an $E' \in \mathcal{B}_s$ with $\mu_s(E'|E) = 1$ such that either $c(w) > c'(w)$ for all $w \in E'$ or $c'(w) > c(w)$ for all $w \in E'$: the reported preferences would thus violate a minimal form of monotonicity as well as the STP.

As indicated in the Introduction, we wish to interpret each of the ϕ_s functions as reflecting the DM's inherent attitude towards actions she does not understand. To make this more tenable, we impose a further assumption on $\succeq$ that has the effect of allowing us to set $\phi_s = \phi_{s'}$ for every $s, s' \in S$. Write $\theta_{s,c}$ for the probability measure on $\mathcal{B}$ defined as $\theta_{s,c}(Z) = \mu_s\{w : c(w) \in Z\}$ (note this is well-defined as derivative actions are $\mathcal{B}$-measurable).

Assumption 2 (State Independence) *If s, s' are non-null, $\theta_{s,c} = \theta_{s',c'}$, and $\theta_{s,c''} = \theta_{s',c'''}$ then:*

$$c \succeq_s c'' \iff c' \succeq_{s'} c'''$$

It is clear that we could obtain a "state dependent" version of Theorem 1 below if we were to drop Assumption 2. We do not pursue this project here.

We now introduce a further choice set. Denote the set of countably additive probability measures on $\mathcal{B}$ using $\Delta(X)$. The set of *Anscombe-Aumann acts*, $\mathcal{F}$, is then the set of all mappings from S to $\Delta(X)$, with typical elements f, f'. Any f is interpreted as a prospect that pays out a lottery with payoff distribution $f(s)$ in the event of any subjective state s, just as in **?**, except with S taking the place of an objective state space. We write $f(s)(E)$ for the probability of $E \in \mathcal{B}$ under measure $f(s)$. The DM's preferences over Anscombe-Aumann acts are given by the relation $\succeq^{AA}$ on $\mathcal{F}$ with $\succ^{AA}$ and $\sim^{AA}$ denoting the respective asymmetric and symmetric components of $\succeq^{AA}$ as usual.

We require $\succeq^{AA}$ to be consistent with the following.

Assumption 3 (AA-EU) *There exists a bounded, continuous, strictly increasing function $v : X \to \mathbb{R}$ and a unique probability measure on $\mathcal{B}$, denoted*

π, such that:

$$f \succeq^{AA} f' \iff \sum_{s \in S} \pi(s)\mathbb{E}_{f(s)}[v(x)] \geq \sum_{s \in S} \pi(s)\mathbb{E}_{f'(s)}[v(x)] \tag{3}$$

for all $f, f' \in \mathcal{F}$.

As with Assumption 1, AA EU can be generated from more basic conditions on $\succeq^{AA}$. These are described in **?**

There is a sense in which the set of derivative actions and the set of Anscombe-Aumann acts intersects, and the next restriction implies that $\succeq^{AA}$ and $\succeq$ are isomorphic over this intersection. To see this, write x for the degenerate lottery that pays out x, and define $\mathcal{F}_\delta$ as the set of Anscombe-Aumann acts such that for every $s \in S$, $f(s) = x$ for some x. Then define $\mathcal{C}_\delta$ as the set of derivative actions that satisfy $c(\{s, g\}) = c(\{s, g'\})$ for all $s \in S$ and any $g, g' \in \mathcal{G}$. Clearly, for any $c \in \mathcal{C}_\delta$ there exists some $f \in \mathcal{F}_\delta$ such that $f(s) = x$ iff $c(s, g) = x$ for all s, and for any $f \in \mathcal{F}_\delta$ there is a $c \in \mathcal{C}_\delta$ with the property $c(s, g) = x$ iff $f(s) = x$ for all s. Where $f \in \mathcal{F}_\delta$, use c_f to refer to the member of $\mathcal{C}_\delta$ that yields the same payoff in each state as f does.

In a similar way, we may also connect the set of actions with the space of derivative actions. For each action a, use c_a for the derivative action that satisfies $c_a(s, g) = g(a)$ for all s and note that c_a pays out x if and only if a turns out to yield x.

Our final restriction connects the three preference relations as follows.

Axiom 2 (Reduction) *The relations $\succeq^*$, $\succeq^{AA}$, and $\succeq$ are mutually consistent insofar as:*

a. For any $a, a' \in \mathcal{A}^*$, $a \succeq^* a'$ if and only if $c_a \succeq c_{a'}$.

b. For any $f, f' \in \mathcal{F}_\delta$, $f \succeq f'$ if and only if $c_f \succeq^{AA} c'_f$.

One way of interpreting Reduction is as the requirement that the DM regards the pairs a and c_a, and c and f_c as identical prospects. This only makes sense if the DM thinks of S as an exhaustive account of what might happen – that is, if $E \subset S$ does not occur, $S \setminus E$ must – and that she knows that all possible consequences of the actions lie within X. Since it is always possible to imagine a catch-all contingency ("none of the above occurs"), requiring the DM to reason in this way does not seem overly demanding (it is implicit in the account of our framework presented in Section 2). And for many economic problems such as investment or policy decisions, it would be taken for granted that the set of consequences is known (for example, X might be a set of monetary quantities, measured in equivalent terms).

We are now ready to state our first representation theorem.

Theorem 1 *The following two claims are equivalent:*

1. $\succeq^*$, $\succeq$, $\succeq^{AA}$ *satisfy Monotonicity, Reduction, Derivative-SEU, State Independence, and AA-EU.*

2. *There exist bounded, continuous, strictly increasing real maps* u *and* ϕ, *and a set of probability measures on each of* $\{\mathcal{B}_s\}_{s\in S}$, $\{\mu_s\}_{s\in S}$ *and on* $\mathcal{E}$, *such that for every* $a, a' \in \mathcal{A}$*:*

$$a \succeq^* a' \qquad \textit{if and only if}$$
$$\sum_{s\in S} \pi(s) u \left(\int_{W_s} \phi(c_a(w))\mathrm{d}\mu_s \right) \geq \sum_{s\in S} \pi(s) u \left(\int_{W_s} \phi(c_a(w)\mathrm{d}\mu_s \right) \quad (4)$$

And, furthermore, π *is unique, the measures* μ_s *are unique whenever* s *is non-null,* ϕ *is unique up to a positive affine transformation, and whenever* $\tilde{\phi} = \alpha\phi + \beta$, *the corresponding* $\tilde{u}$ *satisfies* $\tilde{u}(\alpha x + \beta) = u(x)$ *for all* $x \in \phi(X)$.

4 Ignorance Aversion

In Example 1, we explained the DM's violation of SEU by appeal to a general preference for actions she better understood over those she did not. We called this general tendency ignorance aversion. In this section we provide a formal behavioural definition of what it is for a DM to be ignorance averse and conditions under which one DM may be said to be more ignorance averse than another. We show that, under the representation of Theorem 1, these have neat mathematical characterisations in terms of the convexity of the u function.

To define ignorance aversion formally, we need to introduce some more terminology. Begin by noting that, under the representation, for every $f \in \mathcal{F}$ there is a unique $f' \in \mathcal{F}_\delta$ such that $f'(s) = f'(s')$ for all $s, s' \in S$ and $f' \sim^{AA} f$. In other words, for any $f \in \mathcal{F}$ there is an $f' \in \mathcal{F}$ that yields the same certain payoff in each state and is such that the DM is indifferent between it and f. We refer to $c_{f'}$ – f' relabelled as a derivative action – the *certainty equivalent* of f and denote it $\mathrm{ce}(f)$.

Now define a *ϕ-risk free* derivative action as any $c \in \mathcal{C}$ such that $\int_{W_s} \phi(c(w))\mathrm{d}\mu_s = \int_{W_s} \phi(c(w))\mathrm{d}\mu_{s'}$ for all non-null $s, s' \in S$. That is, c is ϕ-risk free whenever it pays out the same – in terms of the expected value of ϕ – in every subjective state. Of course, a ϕ-risk free derivative action may not be devoid of payoff uncertainty, as it could be that the DM does not understand it and thus considers multiple payoffs possible at various states. We call this sort of uncertainty *unconceptualised uncertainty.*

The intuition behind our definition of ignorance aversion is as follows. Suppose c is a ϕ-risk free derivative action and f is an Anscombe-Aumann act that induces the same probability measure on $\mathcal{B}$ (given π and $\{\mu_s\}_{s\in S}$). f's payoff is uncertain in a way that the DM understands (it depends only on the realisation of the subjective state and the outcome of a lottery), while any uncertainty over c's payoff is unconceptualised. Since each of c and f offers the same distribution over payoffs, a DM who prefers to gamble on uncertainty she understands – that is, an ignorance averse DM – should prefer f to c and hence report $\mathrm{ce}(f) \succeq c$.

To state our defintion, we thus need to characterise the probability measures on $\mathcal{B}$ induced by any f and c given the subjective probabilities in the representation. For any Anscombe-Aumann act, f, define η_f as the probability measure on $\mathcal{B}$ satisfying:

$$\eta_f(E) = \sum_{s\in S} \pi(s) f(s)(E)$$

for all $E \in \mathcal{B}$. Similarly, for any $c \in \mathcal{C}$, let η_c be the probability measure on $\mathcal{B}$ such that:

$$\eta_c(E) = \sum_{s\in S} \mu_s \left(c^{-1}(E) \cap W_s\right)$$

for all $E \in \mathcal{B}$. Observe that since derivative actions are $\mathcal{B}$-measurable, $\eta_c(E)$ is defined for all $E \in \mathcal{B}$.

We may now define ignorance aversion formally as follows.

Definition 1 (Ignorance Aversion) *The DM is* ignorance averse *iff for any ϕ-risk free derivative action c and Anscombe-Aumann act f such that $\eta_f = \eta_c$, $ce(f) \succeq c$.*

In Assumptions 1 and 2 we characterised the DM's attitude to unconceptualised uncertainty using ϕ, while in Assumption 3 we characterised the DM's attitude to uncertainty she does understand – which we require to be the same as her attitude to risk – using the function v. In the proof of Theorem 2 below we show that the two attitudes are related by $v = u \circ \phi$. It is well known that aversion to risk is equivalent to the concavity of the function v and in the Appendix we show that aversion to unconceptualised uncertainty is equivalent to the concavity of ϕ under our representation. It should not, therefore, be surprising that ignorance aversion – which means greater aversion to unconceptualised uncertainty than risk – is equivalent to ϕ being a concave transform of v, that is, to u being convex. This is precisely what our next result establishes.

Proposition 1 *The DM is ignorance averse iff the function u is convex.*

In light of Proposition 1, we may say that u represents the DM's attitude towards ignorance.

Suppose now there are two DMs, A and B, and we wish to compare their attitudes towards ignorance. Denote A's preferences over $\mathcal{C}$ by $\succeq^A$ and B's by $\succeq^B$. If their beliefs and risk preferences are the same and A prefers some ϕ-risk free derivative action c (which she might not understand) to the certainty equivalent of an Anscombe-Aumann act f (which she does understand), then, if she is more ignorance averse than B, B must also prefer c to f. This is the content of the following definition, where we use π^A and μ_s^A to refer to the DM A's beliefs under the representation and v^A for DM A's attitude to risk.

Definition 2 ("More Ignorance Averse") *DM A is more ignorance averse than DM B iff $\pi^A = \pi^B$, $\mu_s^A = \mu_s^B$ for all non-null s, $v^A = v^B$ and, for any ϕ-risk free derivative action c and Anscombe-Aumann action f:*

$$c \succeq^A ce(f) \implies c \succeq^B ce(f)$$

Our next result shows that A's being more ignorance averse than B has a mathematical characterisation that is analogous to that of her being more risk averse in SEU or more ambiguity averse in KMM's representation.

Proposition 2 *Suppose A and B are two DMs whose preferences are represented as in Theorem 1, with u_A and u_B representing their respective attitude towards ignorance. Then A is more ignorance averse than B iff there exists some strictly increasing convex function ψ, such that $u_A = \psi \circ u_B$.*

As noted by KMM when stating a parallel result, Proposition 2 implies that if u_A and u_B are twice continuously differentiable, then A is more ignorance averse than B iff:

$$\frac{u''_A(x)}{u'_A(x)} \geq \frac{u''_B(x)}{u'_B(x)}$$

Thus, provided the differentiability conditions are satisfied, one might refer to $u''_A(x)/u'(x)$ as the *coefficient of absolute ignorance aversion.*

NOTE: check this back against KMM. (1) Do we have an iff? (2) Need some qualification over the domain of x for which this holds.

5 Extension to Source Preference

It has been argued (for example, in **?** and **?**) that DMs' choices between uncertain prospects may hinge on the *source* of uncertainty these prospects' payoffs depend on, where a source may be thought of as a distinct algebra of events. Such decision-making may be irreconcilable with the representation of Theorem 1, as the following example shows.

Example 2: Imagine the oil trader from Example 1 is presented with a choice set that includes c_1 and the action c_3 = "Invest \$1 million on the NASDAQ index, liquidating the position in 6 months' time". For simplicity, suppose her subjective state space is now made up of only the states s_1 and s_2, which concern whether war breaks out in the Middle East, as in Example 1. She thinks that if s_1 occurs, c_3 might yield anything between \$700,000 and \$1.3 million.

She is then offered to choose between the derivative actions c'_1, c''_1, c'_3, and c''_3 below:

c'_1 = "If c_1 pays out more than \$1.1 million and s_1 occurs, receive \$100, otherwise receive \$0"

c''_1 = "If c'_1 pays out \$100, receive \$0, otherwise receive \$100"

c'_3 = "If c_3 pays out more than \$1.2 million and s_1 occurs, receive \$100, otherwise receive \$0"

c''_3 = "If c'_3 pays out \$100, receive \$0, otherwise receive \$100"

The trader reports strict preferences for c'_1 over c'_3 and for c''_1 over c''_3. It is straightforward to verify that this is inconsistent with Monotonicity and Derivative-SEU.

However, the DM may rationalise her preferences as follows. She does not understand any of the members of the choice set and regards the events of c'_1 and c'_3 paying out \$100 as roughly equally likely. But whereas the payoffs from c'_1 and c''_1 depend on a "source of uncertainty" – namely the payoffs resulting from c_1 – about which she, as an oil trader, considers herself an expert, those from c'_3 and c''_3 depend on a source she feels less comfortable speculating on. This is consistent with **?**'s (**?**) "competence hypothesis".

Preferences such as those described in Example 2 may be accommodated in a generalised version of Theorem 1.

To show this, we want to differentiate between sources of uncertainty in terms of actions, so, for any $A \subseteq \mathcal{A}$, let $\mathcal{C}_A \subseteq \mathcal{C}$ be the set of A-*derivatives*, defined as $\{c : g(a) = g'(a) \text{ for all } a \in A \text{ implies } c(\{s,g\}) = c(\{s,g'\})\}$. A-derivatives are derivative actions whose payoff depends only on the true subjective state and the payoff-profile generated by the actions in A. Define $W_{A,s}$ as the finest partition of W_s with the property that $g(a) = g'(a)$ for all $a \in A$ implies $\{s,g\}$ and $\{s,g'\}$ belong to the same element of $W_{A,s}$. Then let $\mathcal{B}_{A,s}$ be the Borel σ-algebra generated by the relative topology on $W_{A,s}$.

Once again we assume that $\succeq$ satisfies Independence so that the preference relation $\succeq_s$ is defined for every $s \in S$. This allows us to define a source as follows:

Definition 3 $\{\mathcal{B}_{A,s}\}_{s\in S}$ *forms a* source *if and only if, it satisfies:*

i. For all non-null s, there exists a bounded, continuous, strictly increasing function $\phi_{A,s} : X \to \mathbb{R}$ and a probability measure on $\mathcal{B}_{A,s}$, denoted $\mu_{A,s}$, such that for all $c, c' \in \mathcal{C}_A$:

$$c \succeq_s c' \iff \int_{W_{A,s}} \phi_{A,s}\left(c(w)\right) \mathrm{d}\mu_{A,s} \geq \int_{W_{A,s}} \phi_{A,s}\left(c'(w)\right) \mathrm{d}\mu_{A,s} \quad (5)$$

and for at least one non-null s, there is a $E \in \mathcal{B}_{A,s}$ such that $1 > \mu_{A,s}(E) > 0$.

ii. There is no $A' \supset A$ such that $\left\{\mathcal{B}_{s,A'}\right\}_{s\in S}$ satisfies part (i).

? give minimal conditions on which a source may be distinguished by the DM's preferences; **?** then gives behavioural axioms under which part (i) of the definition may be satisfied. Abusing terminology, we say action a belongs to source $\{\mathcal{B}_{A,s}\}_{s\in S}$ whenever $a \in A$.

Our generalised version of Theorem 1 weakens Derivative-SEU to the following.

Assumption 4 (Source Dependence) *Every action in $\mathcal{A}$ belongs to a source.*

For any source, we wish to ensure that $\phi_{A,s} = \phi_{A,s'}$ from (5) for all $s, s' \in S$. As before, this will make it possible to talk of the DM's ignorance attitude with respect to a certain source of uncertainty. To achieve this we need to impose a somewhat weaker form of State Independence to that in Section 3.

Assumption 5 (State Independence*) *If s, s' are non-null, c, c', c'', c''' belong to $\mathcal{C}_A$ for some source A, $\theta_{s,c} = \theta_{s',c'}$ and $\theta_{s,c''} = \theta_{s',c'''}$ then:*

$$c \succeq_s c'' \iff c' \succeq_{s'} c'''$$

Once again, a "state dependent" rendering of Theorem 2 below would be possible in the absence of Assumption 5.

A final behavioural condition, which is implied by Assumption 1 but not by Assumption 4, is that the set of all A-derivatives for all any source A is linearly ordered by $\succeq$.

Axiom 3 (Ordering) *Let $\mathbf{A}$ be the set of all sources. $\succeq$ is transitive and complete on $\bigcup_{A\in\mathbf{A}} \mathcal{C}_A$.*

Note that Ordering allows for substantial incompleteness of $\succeq$ over $\mathcal{C}$. If one thinks of the derivative actions whose payoffs depend on the full payoff profile over $\mathcal{A}$ as being the "most complicated" derivative actions in $\mathcal{C}$, Ordering means that the DM only needs to form preferences over the most complicated derivative actions in case there is a source to which every action belongs.

Given Source Dependence, we say the DM *understands* action a if and only if $a \in A$ and, for all non-null s, $\mu_{A,s}(E) = 1$ where $E \subseteq \{\{s, g\} : g(a) = x\}$ for some x. That is, the DM understands a if she believes that S resolves all payoff uncertainty pertaining to a.

Theorem 2 *The following two claims are equivalent:*

1. $\succeq^*$, $\succeq$, $\succeq^{AA}$ *satisfy Reduction, Monotonicity, Ordering, Source Dependence, State Independence*, and AA-EU.*

2. *Every action a belongs to a source $A(a)$; there is a bounded, continuous, strictly increasing real map u and a probability measure π on $\mathcal{E}$; for each $A(a)$ there is a bounded, continuous, strictly increasing map $\phi_{A(a)}$ and a set of probability measures on each of $\left\{\mathcal{B}_{A(a),s}\right\}_{s\in S}$, $\left\{\mu_{A(a),s}\right\}_{s\in S}$; and for any $a, a' \in \mathcal{A}$:*

$$
\begin{aligned}
& a \succeq^* a' \text{ if and only if} \\
& \sum_{s\in S} \pi(s) u_{A(a)} \left(\int_{W_{A(a),s}} \phi_{A(a)} \left(c_a(w)\right) \mathrm{d}\mu_{A(a),s} \right) \\
\geq \; & \sum_{s\in S} \pi(s) u_{A(a')} \left(\int_{W_{A(a'),s}} \phi_{A(a')} \left(c_{a'}(w)\right) \mathrm{d}\mu_{A(a'),s} \right)
\end{aligned}
$$

And, furthermore: π is unique; $\mu_{A(a),s}$ is unique for all $A(a)$ and non-null s; ϕ_A is unique up to an affine transformation, and if $\tilde{\phi}_A = \alpha\phi_A + \beta$, the associated $\tilde{u}_A$ is such that is such that $\tilde{u}_A(\alpha x + \beta) = u_A(x)$ for $x \in \phi(X)$; and for all a, $A(a)$ is unique iff the DM does not understand a, and $a \in A(a')$ for all $a' \in \mathcal{A}$ otherwise.

The uniqueness part of Theorem 2 implies that the set of all actions the DM does not understand may be partitioned according to the source they belong to. Thus, it is only possible for a pair of actions to belong to different sets of sources if the DM understands neither of them.

6 Discussion and Other Literature

TBW

A Appendix

A.1 Proof of Theorem 1

Observe that Derivative-SEU implies that all actions belong to a single source, in which case State Independence* and State Independence are equivalent. Therefore the result follows from Theorem 2. ■

A.2 Lemma A.1

We note the following result (reported as Lemma 6 in KMM), which is invoked in the proofs below.

Lemma A.1 *If $\phi : X \to \mathbb{R}$ is a continuous function and $X \subseteq \mathbb{R}$ is convex, then ϕ is concave iff there exists a $\lambda \in (0,1)$ such that for all $x, y \in X$ where $x \neq y$:*

$$\phi(\lambda x + (1-\lambda)y) \quad \geq \quad \lambda\phi(x) + (1-\lambda)\phi(y)$$

A.3 Proof of Proposition 1

Suppose $\eta_c = \eta_f$ for some ϕ-risk free c and $f \in \mathcal{F}$. Then we have:

$$\sum_{s \in S} \pi(s)\mathbb{E}_{f(s)}[\phi(x)] \quad = \quad \int_{W_{s'}} \phi(c(w))\mathrm{d}\mu_{s'}$$

where s' is non-null. Under the representation, we have:

$$\sum_{s \in S} \pi(s)\mathbb{E}_{f(s)}[u \circ \phi(x)] \quad = \quad u \circ \phi(\mathrm{ce}(f)) \tag{6}$$

If u is convex, then it follows from Jensen's inequality that:

$$\sum_{s \in S} \pi(s)\mathbb{E}_{f(s)}[u \circ \phi(x)] \quad \geq \quad u\left(\int_{W_{s'}} \phi(c(w))\mathrm{d}\mu_{s'}\right)$$

which, under the Representation and given (6), implies that $\mathrm{ce}(f) \succeq c$. Hence, the DM is comparative ignorance averse.

Working in the other direction, let s be such that there exists an $E \in \mathcal{B}_s$ with $0 < \mu_s(E) < 1$. (By Assumption 1, there is such an s). Take any $x, y \in \phi(X)$ and define the derivative action, c, with the property that $c(w) = x$ if $w \in E$, $c(w) = y$ if $w \in W_s \setminus E$, and $c(w) = z$ for all other w where $z := \phi{-}1\left(\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y)\right)$. Now define the A-A act, f, such that $f(s) = l$ where l is a lottery that assigns a probability of $\mu_s(E)$ to consequence x and $(1-\mu_s(E))$ to consequence y, and $f(s')=z$ for all other $s' \in S$.

Observe that c is ϕ-risk free and that $\eta_c = \eta_f$, so if the DM is comparative ignorance averse it must be that $\text{ce}(f) \succeq c$. Using parallel reasoning to that behind (6), it follows that:

$$\sum_{s\in S} \pi(s)\mathbb{E}_{f(s)}[u \circ \phi(x)] \quad \geq \quad \sum_{s\in S} \pi(s) u \left(\int_{W_s} \phi(w(s)) \mathrm{d}\mu_s \right)$$

and hence, by construction, that:

$$\mu_s(E)(u \circ \phi(x))) + (1-\mu_s(E))(u \circ \phi(y)) \geq u\left(\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y)\right) \quad (7)$$

And as (7) holds for any x, y in the convex set $\phi(X)$, Lemma A.1 implies that u is convex. □

A.4 Proof of Proposition 2

For any risk-free derivative action c, write $\text{ce}_A(c)$ for $\phi^{-1}\left(\int_{W_s} \phi(c(w))\mathrm{d}\mu_s\right)$ (where s is non-null). Under the representation we have:

$$c \succeq^A x \iff u_A\left(\phi_A(\text{ce}_A(c))\right) \geq u_A\left(\phi_A(x)\right)$$

Which, since u_A and ϕ_A are strictly increasing, holds iff $\text{ce}_A(c) \geq x$. If $\phi_A = \psi \circ \phi_B$ for a concave ψ then by Jensen's inequality we have $\text{ce}_B(c) \geq \text{ce}_A(c)$ for any risk-free c, from which it follows immediately that A is more ignorance averse than B.

Suppose now that A is more ignorance averse than B and define $\psi := \phi_A \circ \phi_B^{-1}$, which must be strictly increasing under the representation. We proceed as in the proof of Theorem 2 in KMM. Take a non-null s such that there exists a $E \in \mathcal{B}_s$ such that $1 > \mu_s(E) > 0$ and any risk free c. $\text{ce}_B(c) \geq \text{ce}_A(c)$ requires:

$$\phi_B^{-1}\left(\int_{W_s} \phi_B(c(w))\mathrm{d}\mu_s\right) \quad \geq \phi_A^{-1}\left(\int_{W_s} \phi_A(c(w))\mathrm{d}\mu_s\right)$$

which, since $\phi_A = \psi \circ \phi_B$, implies:

$$\psi\left(\int_{W_s} \phi_B(c(w))\mathrm{d}\mu_s\right) \quad \geq \int_{W_s} (\psi \circ \phi_B)(c(w))\mathrm{d}\mu_s \quad (8)$$

For any $x, y \in X$ with $x \neq y$, (8) holds for c such that $c(w) = x$ if $w \in E$ and $c(w) = y$ otherwise. Thus, one can invoke Lemma A.1 to establish that ψ is concave. □

A.5 Proof of Theorem 2

The proof follows a similar path to that for KMM's Theorem 1. We show that the axioms imply the representation and uniqueness properties.

Under Monotonicity and Source Dependence, for source A there is at least one non-null state s such that there exists an $E \in \mathcal{B}_s$ with $\mu(E) \in (0,1)$. By State Independence*, whenever $c, c' \in \mathcal{C}_A$, $c \succeq_s c'$ iff $\mathbb{E}_{\theta_{s,c}}[\phi_{A,s'}(x)] \geq \mathbb{E}_{\theta_{s,c'}}[\phi_{A,s'}(x)]$ for all non-null $s' \in S$. This implies that for any non-null s', s'', $\phi_{A,s'}(x) = \alpha\phi_{A,s''}(x) + \beta$ for some $(\alpha, \beta) \in \mathbb{R}_{++} \times \mathbb{R}$, and hence that for any non-null s' and $c, c' \in \mathcal{C}_A$, $c \succeq_{s'} c'$ iff $\int_{W_{A,s'}} \phi_{A,s}(c(w))\mathrm{d}\mu_{A,s'} \geq \int_{W_{A,s'}} \phi_{A,s}(c'(w))\mathrm{d}\mu_{A,s'}$. Now proceed setting $\phi_A = \phi_{A,s}$

Since ϕ_A is continuous and strictly increasing, for every $c \in \mathcal{C}_A$ and every non-null $s \in S$, there is some unique $x \in X$ such that where $c' \in \mathcal{C}_\delta$ satisfies $c'(\{s,g\}) = x$ for all g, $\{c', s; c\} \sim c$. For each $c \in \mathcal{C}_A$, let c_δ be some member of $\mathcal{C}_\delta$ such that $\{c_\delta, s; c\} \sim c$ for all s. By iterated applications of Independence, for any $c, c' \in \mathcal{C}_A$ it must be that $c \succeq c'$ iff $c_\delta \succeq c'_\delta$.

Reduction then requires that $c \succeq c'$ iff $f_{c_\delta} \succeq f_{c'_\delta}$, which by AA-EU is equivalent to:

$$\sum_S \pi(s)v\left(c_\delta(s)\right) \quad \geq \quad \sum_S \pi(s)v\left(c'_\delta(s)\right)$$

Since ϕ_A and v are both strictly increasing and continuous, there exists some strictly increasing and continuous u_A such that $v = u_A \circ \phi_A$. Hence $c \succeq c'$ iff:

$$\sum_S \pi(s)u_A(\phi_A(c_\delta(s))) \quad \geq \quad \sum_S \pi(s)u_A(\phi_A(c'_\delta(s))) \tag{9}$$

Given Derivative-SEU we have:

$$\{c, s; c_\delta\} \sim_s c_\delta \iff \int_{W_{A,s}} \phi_A(c(w))\mathrm{d}\mu_{A,s} = \phi_A(c_\delta(s))$$

So by construction (9) implies that $c \succeq c'$ iff:

$$\sum_S \pi(s)u_A\left(\int_{W_{A,s}} \phi_A(c(w))\mathrm{d}\mu_{A,s}\right) \quad \geq \quad \sum_S \pi(s)u_A\left(\int_{W_{A,s}} \phi_A(c'(w))\mathrm{d}\mu_{A,s}\right)$$

And then Reduction yields that for any $a, a' \in A$, $a \succeq^* a'$ iff $c_a \succeq c_{a'}$.

Finally, consider any $a, a' \in \mathcal{A}$. Source Dependence implies that there exist sources $A(a)$ and $A(a')$ such that $a \in A(a)$ and $a' \in A(a')$. By Reduction, Ordering, and the reasoning above, it must be that $a \succeq a'$ iff $(c_a)_\delta \succeq (c_{a'})_\delta$ iff $f_{(c_a)_\delta} \succeq f_{(c'_a)_\delta}$. The latter implies:

$$\sum_S \pi(s)v\left((c_a)_\delta(s)\right) \quad \geq \quad \sum_S \pi(s)v\left((c_{a'})_\delta(s)\right)$$

Which as we have shown is equivalent to:

$$\sum_S \pi(s) u_{A(a)} \left(\int_{W_{A(a),s}} \phi_{A(a)}(c(w)) \mathrm{d}\mu_{A(a),s} \right) \geq$$
$$\sum_S \pi(s) u_{A(a')} \left(\int_{W_{A(a'),s}} \phi_{A(a')}(c'(w)) \mathrm{d}\mu_{A(a'),s} \right)$$

as required.

AA-EU implies imply that π is unique and Derivative-SEU implies that $\mu_{A,s}$ is unique for all non-null s; and it is obvious that if s is null, the representation is valid for any arbitrary $\mu_{A,s}$. By assumption, ϕ_A is unique up to a positive affine transformation and $v = u_A \circ \phi_A$, so it is immediate that if $\tilde{\phi}_A = \alpha\phi + \beta$ then the associated $\tilde{u}_A$ satisfies $\tilde{u}_A(\alpha x + \beta) = u_A(x)$ for $x \in \phi(X)$.

Finally, suppose $a \in A(a) \cap A(a')$ where $A(a) \neq A(a')$. We show that this can only be the case where for all non-null s there is some x such that $\mu_s(\{\{s, g\} : g(a) = x\}) = 1$: that is, where the DM understands a. Clearly, if the DM does understand a, then a belongs to all sources, so the uniqueness claim follows from this.

Imagine that $a \in A(a) \cap A(a')$ where $A(a) \neq A(a')$ and that for some s there is an x such that $1 > \mu_s(\{\{s, g\} : g(a) = x\}) > 0$. Definition 3 implies $A(a) \not\subseteq A(a')$ and $A(a) \not\supseteq A(a')$, so there exists an $a'' \in A(a) \backslash A(a')$; since $a'' \notin A(a')$, the DM does not understand a''. The fact that the DM does not understand a implies that whenever $c \succeq_s c'$ iff $\int_{W_{a,s}} \phi_a(c(w)) \mathrm{d}\mu_s \geq \int_{W_{a,s}} \phi(c'(w)) \mathrm{d}\mu_s$ for $c, c' \in \mathcal{C}_a$, ϕ_a is unique up to a positive affine transformation. Since the same holds for a'', it follows that $c \succeq_s c'$ iff $\int_{W_{a,s}} \phi_a(c(w)) \mathrm{d}\mu_s \geq \int_{W_{a,s}} \phi(c'(w)) \mathrm{d}\mu_s$ for all $c, c' \in \mathcal{C}_{a''}$ and hence (given Ordering) that whatever sources a belongs to, a'' also belongs to, a contradiction. ■

References