Centre for
Climate Change
Economics and Policy

The Munich Re Programme: *Evaluating the Economics of Climate Risks and Opportunities in the Insurance Sector*

Grantham Research Institute on
Climate Change and
the Environment

A representation result for choice under conscious unawareness

Oliver Walker and Simon Dietz

September 2011

Centre for Climate Change Economics and Policy
Working Paper No. 68

Munich Re Programme Technical Paper No. 10

Grantham Research Institute on Climate Change and the Environment

Working Paper No. 59

The Grantham Foundation for the Protection of the Environment
E·S·R·C Economic & Social Research Council
Munich RE
University of Leeds
LSE The London School of Economics and Political Science

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council and has five inter-linked research programmes:

1. Developing climate science and economics
2. Climate change governance for a new global deal
3. Adaptation to climate change and human development
4. Governments, markets and climate change mitigation
5. The Munich Re Programme - Evaluating the economics of climate risks and opportunities in the insurance sector (funded by Munich Re)

More information about the Centre for Climate Change Economics and Policy can be found at: http://www.cccep.ac.uk.

The Munich Re Programme is evaluating the economics of climate risks and opportunities in the insurance sector. It is a comprehensive research programme that focuses on the assessment of the risks from climate change and on the appropriate responses, to inform decision-making in the private and public sectors. The programme is exploring, from a risk management perspective, the implications of climate change across the world, in terms of both physical impacts and regulatory responses. The programme draws on both science and economics, particularly in interpreting and applying climate and impact information in decision-making for both the short and long term. The programme is also identifying and developing approaches that enable the financial services industries to support effectively climate change adaptation and mitigation, through for example, providing catastrophe insurance against extreme weather events and innovative financial products for carbon markets. This programme is funded by Munich Re and benefits from research collaborations across the industry and public sectors.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training in climate change and the environment. The Institute is funded by the Grantham Foundation for the Protection of the Environment, and has five research programmes:

1. Use of climate science in decision-making
2. Mitigation of climate change (including the roles of carbon markets and low-carbon technologies)
3. Impacts of, and adaptation to, climate change, and its effects on development
4. Governance of climate change
5. Management of forests and ecosystems

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: http://www.lse.ac.uk/grantham.

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

August 24, 2011

Abstract

There are many examples in policy-making, investment and day-to-day life where the set of contingencies the decision maker can conceive 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. We then specify a choice setting in which the decision maker has preferences over a set of actions stated naturally in English, and over a set of *caveats*. A caveat maps from the set of *permutations* – the product space of the set of contingencies she can conceive of (her *subjective state space*, S) and the set of payoff assignments to the actions – to a space of consequences. We obtain a representation result under which she prefers action a to a' if and only if $\mathbb{E}_\pi u\left[\mathbb{E}_{\mu_s}\phi(w(a))\right] \geq \mathbb{E}_\pi u\left[\mathbb{E}_{\mu_s}\phi(w(a'))\right]$, where π and μ_s are subjective priors on S and each payoff-profile under subjective state s respectively, ϕ and u are utility functions, and $w(a)$ refers to the payoff a yields in assignment w. By endowing the decision maker with beliefs over the set of payoff assignments, we make choice in cases where

*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

conscious unawareness is a major concern (e.g. climate-change policy) tractable by means of some of the standard analytical tools of risk and ambiguity analysis. The representation also allows us to characterise the decision maker's attitude towards perceived payoff uncertainty arising from factors she is unaware of, using properties of the function ϕ; in particular, we say she is *ignorance averse* if and only if ϕ is concave. Using the same framework, we are 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

JEL Classification Numbers: D81, Q54

1 Introduction

Savage's (1954) 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 can conceive of 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 formulate a full state space as in Savage's framework. Rather, it is likely that the set of contingencies the trader can conceive of – her *subjective state space* – omits certain relevant details or distinctions and thus does not resolve all of the trader's payoff uncertainty. Where this is so, we say the trader is *unaware* of the full state space.

Clearly, no DM could ever take her own unawareness into account by considering a list of relevant features that her subjective state space fails to include: if she could do this, she would not be unaware of these features to begin with. She may, nonetheless, reason about the possibility that there exist certain (unspecified) aspects of the full state space of which she is unaware. We describe a DM who believes she is unaware of the full state space as *consciously unaware.*

Savage's theory was evidently not designed for consciously unaware agents and in our view such DMs may be justified in violating SEU. To illustrate this, consider the following example:

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" and assume for simplicity that the trader is aware of only two possible contingencies relevant to action a: either s_1, "peace holds in the Middle East", or s_2, "war breaks out in the Middle East". If s_1 occurs, she thinks the oil price in six months' time could be anything between \$70 and \$115, while if s_2 occurs, she thinks it could be anything between \$85 and \$130.

She may also go to the casino and gamble at a roulette table. She knows the full set of contingencies relevant to any gamble on the roulette wheel is simply the set of pockets on the wheel where the ball might land. Enumerating these pockets $0, 1, \ldots, 36$, suppose her subjective state space is just the product $S := \{s_1, s_2\} \times \{0, 1, \ldots, 36\}$.

The trader recognises that S does not resolve all the 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 equally. For while there is no state in S that resolves the payoff uncertainty relating to a, every state resolves the payoff uncertainty pertaining to gambles on the roulette wheel. We will say that she thus *understands* gambling on the roulette wheel, but not purchasing oil futures.

Now imagine she is asked to consider the following "derivative" action:

$$a' = \begin{array}{c}\text{If } a \text{ pays out more than \$1.1 million} \\ \text{receive \$1,000, otherwise receive \$0}\end{array}$$

She compares a' with gambles on the roulette wheel of the following form:

$$a_n = \begin{array}{c}\text{Receive \$1,000 if the pocket number} \\ \text{is less than } n\text{, otherwise receive \$0}\end{array}$$

She reports a strict preference for a' over a_n for $n = 0, \ldots, 12$, but strictly prefers a_n to a' for $n = 13, \ldots, 36$. If it were the case that the DM's preferences satisfied SEU[1], then, where E stands for the "event" of a paying out more than \$1.1 million and E^c for the "event" of a not paying out more than \$1.1 million, there would exist some subjective probability measure, π, on the power set of $S \times \{E, E^c\}$ representing the DM's beliefs. Abusing notation by writing $\{13, \ldots, 36\}$ for $\{13, \ldots, 36\} \times \{s_1, s_2\} \times \{E, E^c\}$ and E for $E \times S$, under SEU her strict preference for a_{13} over a' would imply:

$$\begin{array}{rl} & \pi(\{0, \ldots, 12\})u(1,000) + (1 - \pi(\{0, \ldots, 12\}))u(0) \\ > & \pi(E)u(1,000) + (1 - \pi(E))u(0) \end{array} \tag{1}$$

[1] We imagine payoffs from all actions are received contemporaneously.

The DM is then asked to consider the "complementary" actions of the form:

$$
\begin{aligned}
a'' &= \text{If } a' \text{ pays out \$1,000 receive \$0, otherwise receive \$1,000} \\
a_n' &= \text{If } a_n \text{ pays out \$1,000 receive \$0, otherwise receive \$1,000}
\end{aligned}
$$

Since it must be that $\pi(\{0,\ldots,12\}) = 1 - \pi(\{13,\ldots,36\})$ and $\pi(E^c) = 1 - \pi(E)$, (1) implies:

$$
\begin{aligned}
& \pi(E^c)u(1,000) + (1-\pi(E^c))u(0) \\
> \; & \pi(\{13,\ldots,36\})u(1,000) + (1-\pi(\{13,\ldots,36\}))u(0)
\end{aligned}
$$

And therefore, under SEU, her preference for a_{13} over a' commits her to strictly preferring a'' to a_{13}'.

The DM might, however, reasonably prefer a_{13}' to a'', explaining her choice as follows. The "event" E, unlike $\{0,\ldots,12\}$, only describes a set of payoffs for a, rather than the set of states of the world in which these payoffs are received, and the DM can conceive of no contingency that corresponds to these payoff outcomes. She therefore does not understand actions a' and a''. By contrast, the DM can conceive of the contingencies that resolve the payoff uncertainty that pertains to a_{13} and a_{13}' and she thus understands these actions. She has a general preference for pursuing actions that she understands – a preference that SEU cannot accommodate – and hence prefers a_{13} to a' and a_{13}' to a''.

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:

$$\sum_{s\in S}\pi(s)u\left(\int_{W_s}\phi(w(a))\mathrm{d}\mu_s\right) \geq \sum_{s\in S}\pi(s)u\left(\int_{W_s}\phi(w(a'))\mathrm{d}\mu_s\right) \quad (2)$$

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 (2) and the way we derive it from our assumptions are familiar from Klibanoff et al.'s (KMM, 2005) 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 (IPCC (2007); Stainforth et al. (2007)), and our understanding of the interface between the climate and the economy is thought to be similarly incomplete (e.g. Heal and Kristroem (2002); Stern (2007); Weitzman (2009)). 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 Mukerji (1997), Ghirardato (2001)) 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.

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 Choice and Representation

The DM chooses from a set of *actions* (denoted $\mathcal{A}$ with typical members written a, a') that are sentences in English describing possible things to do, such as "Build sea defences", "Invest \$5 million in Microsoft", or "Mow the lawn". Her preferences over $\mathcal{A}$ are represented by the binary relation $\succeq^*$ on $\mathcal{A}$ with 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 preferences, first let there be a consequence space X with generic elements x, x', equal to some bounded interval on the real line. One way to interpret X is as encompassing all the *ex ante* monetary valuations the DM might attribute to the result of an action. Such an interpretation would be consistent with investment choices, for example. Second, we assume our DM is endowed with a finite topological 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. The assumptions we make later on will implicitly require that for any $E \in \mathcal{E}$, $S \setminus E$ is conceived of as equivalent to "E does not occur", so we always interpret S as a collectively exhaustive account of what might happen, according to what the DM can conceive of. However, we do not assume that the DM can conceive of every detail relevant to her choice, that is, any s might not be so finely described that the DM would know the payoff any action would lead to were s to occur.

Let $\mathcal{B}$ be the Borel σ-algebra of X and 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 Anscombe and Aumann (1963), except with S taking the place of an objective state space. The DM's preferences over Anscombe-Aumann acts are given by the relation $\succeq^{AA}$ on $\mathcal{F}$. We discuss the observability of $\succeq^{AA}$ in Section 5 below.

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 $g(a)$. (We will usually write $w_s(a)$ to refer to $g(a)$ where this is so.) 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 caveats* 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 *caveats* and written c, c' etc. A caveat is interpreted as a prospect that pays out some amount depending on what permutation obtains: its structure can thus be similar to the "derivative" actions described in Example 1. If the subjective state s occurs and the payoff profile over $\mathcal{A}$ turns out to be g, then caveat 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. Supposing the space of actions is just $\{a, a_0, \ldots, a_{36}\}$, and writing $s_{1,0}$ for the subjective state "Peace holds in the Middle East and the ball lands in pocket number 0", $W_{s_{1,0}}$ is the set of permutations where $s_{1,0}$ obtains. For instance, "$s_{1,0}$ is true, a pays out x, a_0 pays out x_0, …, a_{36} pays out x_{36}" is a member of $W_{s_{1,0}}$. Where $g(a) = x$ and $g(a_i) = x_i$ for $i = 0, \ldots, 36$, we use the compact notation $\{s_{1,0}, g\}$ for this permutation.

a', which we described as a "derivative action" in the Introduction, is really a caveat in this framework, so we write it as c here. Under all permutations where a pays out more than \$1.1 million, it delivers \$1,000; in all others it pays nothing. Thus, 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 asymmetric and symmetric components $\succ$ and $\sim$ (we consider the observability of this relation later on in Section 5).

Our first restriction on $\succeq$ is a familiar independence condition. To state it, we use the notation $\{c, s; c'\}$ to refer to the caveat 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 (Independence) *For any $c, c', c'' \in \mathcal{C}$, $c \succeq \{c', s; c\}$ iff $\{c, s; c''\} \succeq \{c', s; c''\}$.*

Given Independence, 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 Wakker (1985).

Assumption 1 (Caveat-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:*

$$c \succeq_s c' \iff \int_{W_s} \phi_s(c(w))\mathrm{d}\mu_s \geq \int_{W_s} \phi_s(c'(w))\mathrm{d}\mu_s \tag{3}$$

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 our presentation, this is the essence of conscious unawareness. It corresponds to Walker's (2011) characterisation of "the DM believes that if s occurs she may be unaware of something" as her willingness to gamble on some action paying out both more than and less than some payoff x in the event of s.

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 caveats 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.

Turning now to $\succeq^{AA}$, we assume that it is consistent with the following.

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

$$f \succeq^{AA} f' \iff \sum_{s \in S} \pi(s)\mathbb{E}_{f(s)}[v(x)] \geq \sum_{s \in S} \pi(s)\mathbb{E}_{f'(s)}[v(x)] \tag{4}$$

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

As with Assumption 1, AA EU can be generated from more basic conditions on $\succeq^{AA}$ from Wakker (1985).

There is a sense in which the set of caveats 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 caveats 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 $f(s) = c(s, g)$ for all s, and for any $f \in \mathcal{F}_\delta$ there is a $c \in \mathcal{C}_\delta$ with the property $c(s, g) = f(s)$ for all s. Where $c \in \mathcal{C}_\delta$, use f_c to refer to the member of $\mathcal{F}_\delta$ that yields the same payoff in each state as c does.

In a similar way, we may also connect the set of actions with the space of caveats. For each action a, use c_a for the caveat 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 under one of the members of S.

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 $c, c' \in \mathcal{C}_\delta$, $c \succeq c'$ if and only if $f_c \succeq^{AA} f_{c'}$.*

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 conceive of an exhaustive subjective state space does not seem overly demanding[2]. And for many economic problems such as investment or

[2]Walker (2011) gives behavioural conditions under which the DM's preferences reveal that she conceives of the full state space in this way.

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 Independence, Reduction, Caveat 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 $a, a' \in \mathcal{A}$:*

$$a \succeq^* a' \quad \text{if and only if}$$
$$\sum_{s\in S} \pi(s) u \left(\int_{W_s} \phi(c_a(w)) \mathrm{d}\mu_s \right) \geq \sum_{s\in S} \pi(s) u \left(\int_{W_s} \phi(c_a(w) \mathrm{d}\mu_s \right) \quad (5)$$

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 $x \in \phi(X)$.

3 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 concavity of the ϕ function.

To define ignorance aversion formally, we need to introduce some more terminology. Say caveat c is *ϕ-risk-free* iff $\int_{W_s} \phi\left(c(w)\right) \mathrm{d}\mu_s = \int_{W_{s'}} \phi\left(c(w)\right) \mathrm{d}\mu_{s'}$ for all non-null $s, s' \in S$. In other words, c is ϕ-risk free whenever it pays out the same – in terms of the expected value of ϕ – in every state. Of course, a ϕ-risk-free caveat 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.

For any caveat, let η_c be the probability measure on $\mathcal{B}$ defined as:

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

Observe that since caveats are $\mathcal{B}$-measurable, $\eta_c(E)$ is defined for all $E \in \mathcal{B}$. For any probability measure η on $\mathcal{B}$, use $\mathbb{E}_\eta[x]$ for the degenerate caveat that pays out $\mathbb{E}_\eta[x]$ under every permutation.

Consider a DM choosing between two ϕ-risk-free caveats which yield the same expected value of ϕ, one of which she understands and one of which she does not. If her preferences can be represented as in Theorem 1 and she is ignorance averse, she would surely opt for the caveat she understands, since the two caveats are identical in all other respects. This is the intuition behind the following definition.

Definition 1 (Ignorance Aversion) *The DM is ignorance averse iff, for any ϕ-risk-free caveat c, $\mathbb{E}_{\eta_c}[x] \succeq c$.*

Just as, in the standard SEU framework, risk aversion is equivalent to the concavity of the DM's utility function and, in KMM, ambiguity aversion is the same as the concavity of the DM's utility function over second-order lotteries, under our representation ignorance aversion is formally characterised by the concavity of the DM's ϕ function, as the next result shows.

Proposition 1 *Given the representation of Theorem 1, the DM is ignorance averse iff ϕ is concave.*

In the Appendix we give a parallel result that shows that u is concave if and only if the DM is averse to mean-preserving spreads of payoffs across states, that is, just in case he displays what we term *ϕ-risk aversion.*

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 about what the caveats are likely to pay out in each state are the same, and A prefers some ϕ-risk-free caveat c (which she might not understand) to a degenerate action x (which she does understand), then, if she is more ignorance averse than B, B must also prefer c to x. This is the content of the following definition.

Definition 2 ("More Ignorance Averse") *DM A is more ignorance averse than DM B iff they share the same beliefs μ_s for all $s \in S$ and for any ϕ-risk-free caveat c and degenerate caveat, x:*

$$c \succeq^A x \implies c \succeq^B x \tag{6}$$

Once again, we find that under our representation, A's being more ignorance averse than B is analogous mathematically to the properties of comparative risk aversion in SEU and comparative ambiguity aversion in KMM.

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

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

$$-\frac{\phi''_A(x)}{\phi'_A(x)} \quad \geq \quad -\frac{\phi''_B(x)}{\phi'_B(x)}$$

Thus, provided the differentiability conditions are met, one might refer to $-\phi''(x)/\phi'(x)$ as the *coefficient of absolute ignorance aversion.*

4 Application to Source Preference

It has been argued (for example, in Heath and Tversky (1991) and Chew et al. (2008)) 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 a and the action b = "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, b might yield anything between \$700,000 and \$1.3 million.

She is then offered to choose between the caveats a_{s_1}, a'_{s_1}, b_{s_1}, and b'_{s_1} below:

a_{s_1} = "If a pays out more than \$1.1 million and s_1 occurs, receive \$1000, otherwise receive \$0"

a'_{s_1} = "If a_{s_1} pays out \$1000, receive \$0, otherwise receive \$1000"

b_{s_1} = "If b pays out more than \$1.2 million and s_1 occurs, receive \$1000, otherwise receive \$0"

b'_{s_1} = "If b_{s_1} pays out \$1000, receive \$0, otherwise receive \$1000"

The trader reports strict preferences for a_{s_1} over b_{s_1} and for a'_{s_1} over b'_{s_1}. Supposing her preferences satisfy Independence, this implies that she violates the Caveat-SEU, because where $E_a = \{\{s, g\} : s = s_1 \text{ and } g(a) > 1,100,000\}$ and $E_b = \{\{s, g\} : s = s_1 \text{ and } g(b) > 1,200,000\}$, $a_{s_1} \succ_{s_1} b_{s_1}$ implies:

$$\begin{aligned} & \mu_{s_1}(E_a)\phi(1000) + (1 - \mu_{s_1}(E_a))\,\phi(0) \\ > \; & \mu_{s_1}(E_b)\phi(1000) + (1 - \mu_{s_1}(E_b))\,\phi(0) \end{aligned}$$

while $a'_{s_1} \succ_{s_1} b'_{s_1}$ implies:

$$\begin{aligned} & (1 - \mu_{s_1}(E_a))\,\phi(1000) + \mu_{s_1}(E_a)\phi(0) \\ > \; & (1 - \mu_{s_1}(E_b))\,\phi(1000) + \mu_{s_1}(E_b)\phi(0) \end{aligned}$$

which is impossible if μ_{s_1} is a probability.

However, the DM may rationalise her preferences as follows. She does not understand either action and regards E_a and E_b as roughly equally likely. But whereas the payoffs from a_{s_1} and a'_{s_1} depend on a "source of uncertainty" – namely the payoffs resulting from a – about which she, as an oil trader, considers herself an expert, those from b_{s_1} and b'_{s_1} depend on a source she feels less comfortable speculating on. This is consistent with Heath and Tversky's (1991) "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-*caveats*, defined as $\{c : g(a) = g'(a) \text{ for all } a \in A \text{ implies } c(\{s, g\}) = c(\{s, g'\})\}$. A-caveats are caveats 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 $c, c' \in \mathcal{C}_A$:

$$c \succeq_s c' \iff \int_{W_{A,s}} \phi_{A,s}\,(c(w))\,\mathrm{d}\mu_{A,s} \geq \int_{W_{A,s}} \phi_{A,s}\,(c'(w))\,\mathrm{d}\mu_{A,s} \quad (7)$$

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).*

Chew and Sagi (2008) give minimal conditions on which a source may be distinguished by the DM's preferences; Wakker (1985) 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 Caveat-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 (7) 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 2.

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-caveats 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 caveats whose payoffs depend on the full payoff profile over $\mathcal{A}$ as being the "most complicated" caveats in $\mathcal{C}$, Ordering means that the DM only needs to form preferences over the most complicated caveats 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, Independence, 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 $a, a' \in \mathcal{A}$:*

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

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 two actions to belong to separate sources if the DM understands neither of them.

5 Discussion

Theorems 1 and 2 show how one can represent mathematically the behaviour of consciously unaware decision makers who violate SEU but observe certain weaker regularity conditions. The results therefore offer foundations for incorporating this kind of decision-making into a wide range of areas of economic theory, including game theory, finance, and policy analysis, as well as facilitating the empirical analysis of such DMs. The fact that these representations maintain much of the technical and intuitive apparatus of SEU means that (we hope) such an integration could be achieved without departing radically from existing analytical methods in these areas.

In this concluding section we discuss two broad issues concerning our representations. The first of these is the observability of the preference relations $\succeq$ and $\succeq^{AA}$, which bear most of the weight of our assumptions. For whereas $\succeq^*$, the DM's preferences over actions, may be elicited straightforwardly, it might be objected that $\succeq$ and $\succeq^{AA}$ are defined on inherently *subjective* objects – namely, maps that are defined in terms of S – and are thus not readily observable.

Walker (2011) shows how, in a choice setting similar to that used here, a DM's subjective state space may be revealed by his preferences over a set of prospects defined in terms of the set $\mathcal{A}$. In principle, one could therefore uncover S using that approach and then test whether the DM's preferences over caveats and Anscombe-Aumann acts were consistent with our assumptions. Though we acknowledge this testing procedure is likely to be infeasible in most practical settings, we note that most other representations in the literature, including SEU and KMM's model, are guilty of the same charge.

The second issue is the relation between the effect of conscious unawareness and ambiguity, in particular the connection between KMM's representation and our own. Our explanation for why the DM in our initial example violated SEU – that she was ignorance averse – closely parallels the explanation of the Ellsberg paradox in terms of ambiguity aversion. Under both accounts, the DM exhibits a tendency to choose prospects whose payoff structure (be it the probability distribution over states or the mapping from states to consequences) is known to her, and this tendency is inconsistent with SEU because it leads to a violation of Savage's P2, the Sure Thing Principle.

To see precisely how each violation of P2 arises, recall Ellsberg's (1961) "single-urn" thought experiment, where a DM may gamble on whether a ball drawn from an urn containing 90 balls is red, green, or yellow. She knows that there are 30 red balls and that each of the remaining 60 balls is either green or yellow, though in unknown proportion. Writing R, G, and Y for the respective events corresponding to "The ball drawn is red / black / yellow", suppose the DM is offered the choice between prospect a, which pays out \$100 in case of R and \$0 otherwise, and prospect b, which pays out \$100 in case of B and \$0 otherwise. Since a and b pay out the same under event Y, the Sure Thing Principle requires that the DM's preference between a and b is independent of the amount they pay under Y (provided it is the same). Thus, if she prefers a to b, she should also prefer a', which pays out \$100 in case of $R \cup Y$ and \$0 otherwise, to b', which pays out \$100 if $G \cup Y$ and \$0 otherwise. In KMM, an ambiguity averse DM – reasoning that the prospects a and b' are ambiguous whereas a' and b are not – may, however, prefer a to b and b' to a'.

In our theory, an analogous case would be where S comprised s_1 and s_2 and W_{s_2} could be paritioned into $\{E, E'\}$ such that $1 > \mu_{s_2}(E) > 0$. The DM could then be offered the choice between caveat c_1, defined such that $c_1(w) = 100$ if $w = \{s_1, g\}$ for any g and $c_1(w) = \$0$ otherwise, and c_2, which satisfies $c_2(w) = \$100$ if $w \in E$ and $c_2(w) = \$0$ otherwise. Under the Sure Thing Principle, the DM prefers c_1 to c_2 if and only if she prefers c_1', which pays out \$100 if $c_1(w) = \$100$ or any permutation in E' occurs and \$0 otherwise, over c_2', which pays out \$100 if s_2 obtains and \$0 otherwise. However, if she is ignorance averse, the DM may reason that she understands

c_1 and c_2' but not c_1' and c_2, and thus report a preference for c_1 over c_2 and c_2' over c_1'.

Although these two cases have very similar structures, the way they are respectively accommodated in KMM's representation and ours is somewhat different. In KMM's treatment of Ellsberg's single-urn example, the DM's preferences reveal that she considers various probability measures on $2^{\{R,G,Y\}}$ possible. All of these measures assign the same probability to R and $G \cup Y$, but not all assign the same probability to G and $R \cup Y$. This is what explains her preferences for a over b and b' over a'. In our treatment of the analogous case, the DM's preferences reveal that she considers multiple payoff profiles under s_2 possible: under all these profiles, c_1 and c_2' yield a determinate payoff, but under different profiles c_2 and c_1' pay out different amounts. Thus, while the violation of Savage's P2 in our theory has uncertainty over the way actions map from states to consequences as its genesis, in KMM the axiom fails because of uncertainty over the true probability measure on events. To cope with examples where both these forms of uncertainty were present, one would need a more general theory that nested both KMM's representation and our own.

A further distinction between the way ambiguity features in KMM and the way conscious unawareness works here is in how the DM's attitude towards them relates to her appetite for risk. In KMM's representation, the DM is portrayed as evaluating prospects, first by computing their expected utility under each possible probability measure (and thus accounting for their riskiness under each measure), and then by calculating their expected value given the likelihood she attaches to each measure (thereby accounting for their ambiguity). Her attitudes towards uncertainty may thus be decomposed into separate attitudes towards risk and ambiguity. Here, by contrast, the DM's attitude towards risk is given by the function v and her attitude to ignorance by ϕ, and the two are connected by the identity $v = u \circ \phi$. Her ignorance attitude is thus a *component* of her risk attitude in our theory.

It should be stressed that this does not mean conscious unawareness could not be used – as KMM's approach has been – to explain phenomena that do not seem reconcilable with SEU and standard measures of risk aversion. An example of this is the "Equity Premium Puzzle" in financial economics, which Collard et al. (2011) have explained in terms of ambiguity aversion. This is simply because it need not be the case under our representation that $a \succeq^* a'$ iff $\mathbb{E}_{\eta_{c_a}}[v(x)] \geq \mathbb{E}_{\eta_{c_{a'}}}[v(x)]$.

Finally, although there is no dynamic aspect to our representation, we note that a dynamic treatment of conscious unawareness might involve departures from the Bayesian paradigm that would not be called for in the presence of ambiguity. Essentially this is because, in the presence of conscious unawareness, the DM may increase her knowledge of the full state space over time, something that is impossible in the presence of ambiguity. Updating in the presence of unawareness is an active area of research and has been studied by Karni and Viero (2011).

A Appendix

A.1 Proof of Theorem 1

Observe that Caveat-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

First note that where c is a risk-free caveat and s' is a non-null state, $\mathbb{E}_{\eta_c}[\phi(x)] = \sum_S \pi(s) \int_{W_s} \phi(c(w))\mathrm{d}\mu_s = \int_{W_{s'}} \phi(c(w))\mathrm{d}\mu_{s'} \cdot \sum_S \pi(s)$. Thus:

$$\mathbb{E}_{\eta_c}[\phi(x)] \quad = \quad \int_{W_s} \phi(c(w))\mathrm{d}\mu_s \tag{8}$$

for all non-null $s \in S$.

Under the representation, $\delta_{\mathbb{E}_{\eta_c}[x]}$ is preferred to the risk-free caveat c iff:

$$\sum_S \pi(s)u\left(\phi(\mathbb{E}_{\eta_c}[x])\right) \quad \geq \quad \sum_S \pi(s)u\left(\int_{W_s} \phi\left(c(w)\right)\mathrm{d}\mu_s\right)$$

which, given (8), holds iff:

$$u\left(\phi(\mathbb{E}_{\eta_c}[x])\right) \quad \geq \quad u\left(\mathbb{E}_{\eta_c}[\phi(x)]\right)$$

Since u is strictly increasing, this is equivalent to:

$$\phi\left(\mathbb{E}_{\eta_c}[x]\right) \quad \geq \quad \mathbb{E}_{\eta_c}[\phi(x)] \tag{9}$$

Now if ϕ is concave, it is immediate from Jensen's inequality that (9) will be satisfied for any risk-free caveat, so the DM is ignorance averse.

Working in the other direction, suppose that the DM is ignorance averse and recall that, for some s there exists a member of $\mathcal{B}_s$, E, such that $1 > \mu(E) > 0$. For any $x, y \in X$ with $x \neq y$, there is a caveat c satisfying $c(w) = x$ for all $w \in E$, $c(w) = y$ for all $w \in E'$, and $c(w) = \phi^{-1}\left(\mu(E)\phi(x) + \mu(E')\phi(y)\right)$ for all $w \in W_{s'}$ where $s' \neq s$. Evidently, c is a risk-free caveat.

The next step is to show that $\mathbb{E}_{\eta_c}[x] \leq \mu_s(E)x + (1-\mu_s(E))y$. Suppose instead that $\mathbb{E}_{\eta_c}[x] > \mu_s(E)x + (1-\mu_s(E))y$. This could only hold if:

$$\phi^{-1}\left(\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y)\right) \quad > \quad \mu_s(E)x + (1-\mu_s(E))y$$

which (since s is non-null) implies:

$$\begin{gathered}\phi^{-1}\left(\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y)\right) > \\ \pi(s)\left(\mu_s(E)x + (1-\mu_s(E))y\right) + \\ (1-\pi(s))\phi^{-1}\left(\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y)\right)\end{gathered}$$

It follows that:

$$\begin{gathered}\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y) > \\ \phi\left[\begin{array}{c}\pi(s)\left(\mu_s(E)x + (1-\mu_s(E))y\right) + \\ (1-\pi(s))\phi^{-1}\left(\mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y)\right)\end{array}\right]\end{gathered} \tag{10}$$

But under the representation, (10) can be satisfied iff $c \succ_s \delta_{\mathbb{E}_{\eta_c}[x]}$. Since by construction it must be that $c \sim_{s'} \delta_{\mathbb{E}_{\eta_c}[x]}$ for all other s', it follows by Independence that $c \succ \delta_{\mathbb{E}_{\eta_c}[x]}$, contradicting the initial supposition that the DM was ignorance averse. We thus conclude $\mathbb{E}_{\eta_c}[x] \leq \mu_s(E)x + (1-\mu_s(E))y$.

If $\mathbb{E}_{\eta_c}[x] \leq \mu_s(E)x + (1-\mu_s(E))y$ then the fact that $\delta_{\mathbb{E}_{\eta_c}[x]} \succeq c$ implies $\delta_{\mathbb{E}_{\mu_s}[x]} \succeq c$ where $\delta_{\mathbb{E}_{\mu_s}[x]}(w) = \mu_s(E)x + (1-\mu_s(E))y$ for all w. Since c is risk-free, this can only hold if $\delta_{\mathbb{E}_{\mu_s}[x]} \succeq_{s'} c$ for all s'. For state s it must therefore be that:

$$\phi\left(\mu_s(E)x + (1-\mu_s(E)y)\right) \quad \geq \quad \mu_s(E)\phi(x) + (1-\mu_s(E))\phi(y) \tag{11}$$

This argument shows that (11) is true for any any x, y where $x \neq y$. Since $\mu_s(E) \in (0,1)$, Lemma A.1 implies that ϕ is concave. □

A.4 Proof of Proposition 2

For any risk-free caveat c, write $\text{ce}_A(c)$ for $\phi^{-1}\left(\int_{W_s} \phi(c(w))\text{d}\mu_s\right)$ (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:

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

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

$$\psi\left(\int_{W_s} \phi_B(c(w))\mathrm{d}\mu_s\right) \geq \int_{W_s} (\psi \circ \phi_B)(c(w))\mathrm{d}\mu_s \tag{12}$$

For any $x, y \in X$ with $x \neq y$, (12) 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 Independence 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:

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

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:

$$\sum_S \pi(s) u_A(\phi_A(c_\delta(s))) \quad \geq \quad \sum_S \pi(s) u_A(\phi_A(c'_\delta(s))) \tag{13}$$

Given Caveat-SEU we have:

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

So by construction (13) 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:

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

as required.

AA-EU implies imply that π is unique and caveat-EU 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 $\mu_s\left(\{\{s, g\} : g(a) = x\}\right) = 1$: 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. ■

A.6 ϕ-Risk Aversion

For each c, write $\phi_c = \sum_S \pi(s) \int_{W_s} \phi(c(w))\mathrm{d}\mu_s$, that is, the expected value of $\phi(x)$ under the probability measures π and $\{\mu_s\}_{s\in S}$. $\phi^{-1}(\phi_c)$ is then the consequence that yields this value of ϕ.

Definition 4 (ϕ-Risk Aversion) *The DM is ϕ-risk averse iff for any $c \in \mathcal{C}$, $\phi^{-1}(\phi_c) \succeq c$.*

Proposition 3 *Under the representation, provided there are at least two non-null states in s, the DM is ϕ-risk averse iff u is concave.*

Proof: This result echoes part of KMM's Proposition 1 and the proof has the same structure.

The DM's being ϕ-risk averse is equivalent to:

$$u\left(\sum_S \pi(s) \int_{W_s} \phi(c(w))\mathrm{d}\mu_s\right) \quad \geq \quad \sum_S \pi(s) u\left(\int_{W_s} \phi(c(w))\mathrm{d}\mu_s\right)$$

for all $c \in \mathcal{C}$. Jensen's inequality thus ensures that if u is concave, the DM is ϕ-risk averse.

Now suppose that the states s and s' are non-null and that the DM is ϕ-risk averse. Note that since ϕ is continuous, $\phi(X)$ – the range of u – is a convex subset of the real line. Take any $x, y \in \phi(X)$ where $x \neq y$ and, letting $a = \phi^{-1}(x)$ and $b = \phi^{-1}(y)$, consider the caveat, c, that satisfies $c(w) = a$ for all $w \in W_s$ and $c(w) = b$ otherwise. Since it must be that $\phi^{-1}(\phi_c) \succeq c$, we have:

$$u(\pi(s)x + (1 - \pi(s))y) \quad \geq \quad \pi(s)u(x) + (1 - \pi(s))u(y)$$

Since $\pi(s) \in (0, 1)$ and given Lemma A.1, this means that u is concave. □

References

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

CHEW, S. H., K. K. LI, R. CHARK, AND S. ZHONG (2008): "Source Preference and Ambiguity Aversion: Models and Evidence from Behavioral and Neuroimaging Experiments," *mimeo: Hong Kong University of Science and Technology.*

CHEW, S. H. AND J. S. SAGI (2008): "Small Worlds: Modeling Attutitudes towards Sources of Uncertainty," *Journal of Economic Theory*, 139, 1–24.

COLLARD, F., S. MUKERJI, K. SHEPPARD, AND J.-M. TALLON (2011): "Ambiguity and the Historical Equity Premium," *mimeo: University of Oxford.*

ELLSBERG, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643–669.

GHIRARDATO, P. (2001): "Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty," *Economic Theory*, 17, 247–276.

HEAL, G. AND B. KRISTROEM (2002): "Uncertainty and climate change," *Environmental and Resource Economics*, 22, 3–39.

HEATH, C. AND A. TVERSKY (1991): "Preference and Belief: Ambiguity and Competence in Choice under Uncertainty," *Journal of Risk and Uncertainty*, 4, 5–28.

IPCC (2007): "Summary for Policymakers," in *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. Tignor, and H. Miller, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

KARNI, E. AND M.-L. VIERO (2011): "Reverse Bayesianism: a Choice-Based Theory of Growing Awareness," *mimeo: John Hopkins University.*

KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): "A Smooth Model of Decision Making under Uncertainty," *Econometrica*, 73, 1849–1892.

MUKERJI, S. (1997): "Understanding the nonadditive probability decision model," *Economic Theory*, 9, 23–46.

SAVAGE, L. J. (1954): *The Foundations of Statistics*, Dover.

Stainforth, D., M. Allen, E. Tredger, and L. Smith (2007): "Confidence, uncertainty and decision-support relevance in climate predictions," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365, 2145.

Stern, N. H. (2007): *The Economics of Climate Change*, Cambridge.

Wakker, P. P. (1985): "Continuous Expected Utility on Arbitrary State Spaces," *Methods of Operations Research*, 50, 113–129.

Walker, O. J. (2011): "Reasoning about Unawareness: A Decision Theoretic Account," *mimeo: University of Oxford.*

Weitzman, M. (2009): "On modeling and interpreting the economics of catastrophic climate change," *The Review of Economics and Statistics*, 91, 1–19.