Economics
Working Papers

2022-06

Lost in objective translation: Awareness of unawareness when unknowns are not simply unknowns

Marie-Louise Vierø

DEPARTMENT OF ECONOMICS AND BUSINESS ECONOMICS
AARHUS UNIVERSITY

Lost in objective translation: Awareness of unawareness when unknowns are not simply unknowns

Marie-Louise Vierø*
Aarhus University

November 6, 2023

Abstract

This paper models decision makers who are cognisant of their own potential unawareness and who may have an a priori sense that not all unknowns are equal. Different types of surprises may be anticipated in different situations, and this matters for choices. The paper proposes a model that accommodates such a priori differences in anticipation. It allows for subjectively different unknowns and provides a representation of preferences that has an expected utility structure, but where the attitudes towards differently perceived unknowns are allowed to differ.

JEL classification: D8; D81; D83
Keywords: Awareness; unawareness; unknown unknowns; utility of indescribable consequences; attitudes towards the unknown.

*Department of Economics and Business Economics, Aarhus University, Fuglesangs Alle 4, 8210 Aarhus V, Denmark. Email: mlv@econ.au.dk. I gratefully acknowledge financial support from Aarhus University Research Foundation. I thank Edi Karni, Burkhard Schipper, seminar participants at Universita Bocconi and University of Graz, and conference participants at D-TEA 2023, DAEiNA 2023, Decision Theory with Unawareness, Aarhus Workshop on Known and Unknown Unknowns, Advances in collective and individual decision-making and information economics, Awareness Workshop, and the Philosophy of Science Association meeting 2021 for comments and suggestions. Any errors are my own.

1 Introduction

This paper models decision makers who are cognisant of their own potential unawareness and who may have an a priori sense that not all unknowns are equal. In particular, a decision maker's a priori anticipation of how potential unknown consequences of her actions will affect her well-being may differ for different actions. In other words, there is a subjective perception that unknowns are not simply unknowns.

By the nature of unawareness, a decision maker who is aware of her unawareness cannot know exactly what she is unaware of. Instead, she is simply aware that there may be aspects of the universe that she cannot describe with her current language. There may be potential consequences of her actions that have never been heard of or imagined.

Karni and Vierø (2017) provides a choice theoretic model of growing awareness when the decision maker has awareness of her unawareness. In that paper, all potential unknown consequences arising from any act are "collapsed" in the same way to a single unknown, which is assigned a utility value that captures the decision maker's expected value of the unknown and is the same for all acts. In other words, all potential consequences that the decision maker is unaware of are assigned the same utility value a priori. As a result, all unknown consequences are treated identically from an a priori point of view.

However, contrary to the assumption in Karni and Vierø (2017), it is reasonable for a decision maker to have a priori differences in the perceptions of unknowns arising in different situations. For example, the decision maker may perceive there to be positive potential unknowns from introducing artificial intelligence into radiology tasks and negative potential unknowns from introducing artificial intelligence into warfare.

A different example is the perception of positive and negative side-effects of medication. Recently, the anticipation of negative unknown, never seen or heard of, side effects of COVID-19 vaccines received a lot of attention and policy interest, as vaccine sceptics raised concerns about negative unknowns. Discoveries of positive unknown side effects play an important role in pharmaceutical research and development. As an example, consider the development of the drug Sildenafil, which was undergoing clinical trials as a treatment for pulmonary hypertension.

The discovery of an unanticipated side effect lead to the drug being developed for what is now its main use. Sildenafil is the active ingredient in Viagra, and judging by its commercial success, the side effect can be classified as positive.[1]

The point of the examples is to illustrate that not only does awareness of unawareness affect choices, so does a decision maker's anticipation of discovering different new consequences depending on her actions. Hence, all unknowns are not equal. One can also think of it as if there are different shades of the unknown. Since the decision maker's different subjective perceptions presumably matter for her choices, a good model of behavior should accommodate this.

The present paper proposes a model of awareness of unawareness that indeed accommodates such a priori differences in anticipation. It allows the decision maker to operate with multiple subjectively different unknowns, and her a priori anticipation of how potential unknown consequences of her actions will affect her well-being may differ for different actions. In particular, it will be assumed that the decision maker perceives some subjective set of potential unknown consequences. These could for example be "something good," "something bad," something terrible," and "a minor inconvenience." This introduces subjectivity into the description of consequences, since they incorporate how the decision maker feels about the unknown. Thus, there are aspects of the decision environment that are not objectively defined.

On the other hand, outside observers can only meaningfully describe an act using objective language, because the only resolutions of uncertainty that we can hope to verify are those that are objectively described. Bets, trades, and contracts must be specified in objective language in order to be settled once uncertainty resolves. A bet can, without causing later problems, specify the delivery of a consequence, which is none of the prior known consequences, if such a consequence is discovered.[2] The same is not the case for subjective consequences. It cannot be observed or verified how a decision maker perceives a new consequence. In other words, while we can observe and verify that some unanticipated consequence materialized, we cannot observe or verify neither how the decision maker anticipated it ex ante, nor how he feels about it ex post. Further complicating the matter,

[1]Pulmonary hypertension is high blood pressure in the blood vessels that supply the lungs. The example is adapted from episode 12 of the Signal podcast with the head of R&D at Pfizer.

[2]This means that the unknown consequence in Karni and Vierø (2017) is verifiable ex post.

different decision makers may disagree on the subjective nature of a consequence.

The lack of verification possibilities makes the question of how many unknowns to include when modelling the decision maker far from trivial. However, it is an important question. It will be argued below that a serious problem arises if we fail to recognise that the decision maker perceives multiple subjective unknowns. Rational decision makers will display seemingly irrational behaviour if we fail to recognise the issue. Thus, the question of whether we should include multiple subjective unknowns needs to be addressed.

Allowing for subjectivity in the description of unknown consequences provides a challenge for the revealed preference approach. In order for the revealed preference approach to be valid, a model needs to be framed in objective language such that the resolution of uncertainty is objectively verifiable. Axioms can only meaningfully be imposed on entities that can be objectively described.

Viewed simply as courses of action, acts are objective entities and can be taken as primitives. However, without more structure, a model is not particularly useful. The issue of subjectivity arises if we want to model uncertainty with a state space and acts as being state-contingent. Then there is a mismatch between the objects the decision maker forms preferences over and the objects we as outside observers can perform meaningful axiomatic choice tests between. While the decision maker can split the atom "unknown consequence" into multiple subjective unknowns, choice tests cannot be performed at this subatomic level. An outside observer's description of the objects of choice, which has to rely on objective language, will in general leave some unresolved residual subjective uncertainty.

The question therefore arises whether we can model a decision maker who has preferences over courses of action that have subjective unknown elements as if he is maximising expected utility over the translations of these to objective state-contingent acts. The present paper provides conditions under which the answer to this question is affirmative.

The results enable us to fit the problem with multiple subjective unknowns into a framework that is as close as possible to existing frameworks. The model is tractable and easy to use in applications. This further motivates why it is desirable to translate the problem to one where the objects of choice are modelled with objective state-contingent acts.

The present paper combines two different approaches. Skiadas (1997a,b) considers preferences over acts, where an act is simply a label of a course of action. This provides a useful way of thinking of acts when there is subjectivity in the description of the unknown consequences. The present paper combines the approach of Skiadas with the approach from the reverse Bayesianism papers by Karni and Vierø (2013, 2015, 2017), Vierø (2021), Karni, Valenzuela-Stookey and Vierø (2021), and Dominiak and Tserenjigmid (2018, 2022).[3]

Awareness of unawareness has been addressed epistemically by Board and Chung (2011), Walker (2014), and Halpern and Rego (2009, 2013). Choice theoretic papers on awareness of unawareness include Walker and Dietz (2011), Alon (2015), Piermont (2017), and Kochov (2018). For an overview of the earlier unawareness literature, see Schipper (2015).

Grant and Quiggin (2015) presents a model where the decision maker may face a favourable or unfavourable surprise. However, they treat the type of surprise as an objective property and therefore bypass most of the issues the present paper addresses. Grant and Quiggin (2013a, 2013b) consider dynamic games with differential awareness, where players may be unaware of some histories of the game. They provide logical foundations for players using inductive reasoning to conclude that there may be parts of the game tree of which they are unaware. Ozbay (2007) considers games where a player's awareness may increase due to strategic announcements by his opponent.

Kreps (1979, 1992) and Dekel, Lipman and Rustichini (2001) derive a subjective state space from a decision maker's choices over menus that constitute future choice sets. Each subjective state corresponds to a possible future preference relation. In Dekel et al., the decision maker anticipates that different future circumstances may exist in which he will have different preferences, and the optimal menu is the one that gives the decision maker the best possibilities ex post, not yet knowing which of his ex post preferences will apply. In contrast, in the present paper the decision maker has to make a final choice right away. She recognizes that she does not have an objective basis for evaluating yet unknown outcomes that potentially arise. Instead she uses her subjective evaluation, or estimate, of the utility she will get from the potential unknown consequences arising from different acts to guide her single final decision. These subjective evaluations need not hold ex post.

[3]Schipper (2022) connects reverse Bayesian updating of beliefs to the statistical literature on exchangeable random partitions.

Blume, Easley and Halpern (2021) considers a decision maker who choses between syntactic programs. The syntactic programs are "if ... then ... else ..." statements that involve tests formulated as propositions about the world. The goal of Blume et al. is to obtain an expected utility representation without taking states and consequences as part of the description of the problem. In the present paper, I use the structure provided by the actions and consequences the decision maker is aware of. This has bite in terms of identification. Given knowledge of actions and consequences, the decision maker's and the outside observer's objective formulation of the possible resolutions of uncertainty coincide. Subjectivity is about attributes that may be revealed later so that the decision maker does not yet know what they are at the time of decision.

Lipman (1999) models a decision maker who lacks logical omniscience. There may be some logically equivalent information sets that the agent does not perceive as such. The state space is extended to include "impossible possible worlds" in which such sets indeed differ. The set of consequences is given. In the present paper, the decision maker acts logically consistent given his subjective ex ante perception of consequences. Because there are multiple subjective perceptions, objective states are not complete descriptions of the world. Objectively verifiable language is insufficient to describe all logical consistencies that affect preferences.

The paper is organised as follows: Section 2 presents the primitives and develops the framework. Section 3 illustrates the problem that arises if we fail to recognise that a decision maker perceives multiple subjective unknowns. Section 4 presents the preference structure and the main representation result. Section 5 provides a test to reveal multiple subjective unknowns, and Section 6 concludes. Proofs are collected in the Appendix.

2 Model

Consider the following framework of a decision maker facing a potential expansion of his awareness. The primitives of the model are a set of basic actions, to be interpreted as names of possible courses of action, and a set of known consequences that the decision maker is aware of. Let A denote the set of basic actions, with generic element a. Assume that A is non-empty and finite. Actions eventually lead to consequences. Let C denote the set of known consequences that is, the set of consequences the decision maker is aware of, also assumed non-empty and finite. The sets A and C are assumed to be objective and observable by outsiders. Thus, everyone agrees on the state of knowledge.

The decision maker is aware of her potential unawareness. She therefore entertains the possibility that in addition to the known consequences in C, new, not yet seen or heard of, and thus not experienced, consequences may result from the actions. Such a new consequence represents an expansion of the decision maker's awareness.

Define $x \equiv \neg C$ to be the abstract "consequence" that has the interpretation "none of the above." The abstract consequence x will be referred to as the objectively unknown, because it can be verified whether or not a new consequence that does not belong to C was discovered. Define $\widehat{C} \equiv C \cup \{x\}$, and denote it the set of *objectively* extended consequences.

2.1 Objective modeling of uncertainty

Given awareness of A and C, the possible resolutions of uncertainty that can be objectively described are given by the objective conceivable state space, defined in Karni and Vierø (2017).[4] It is defined by

$$S_0 \equiv (\widehat{C})^A = \{s : A \to \widehat{C}\}, \tag{1}$$

that is, the set of all functions from the set of basic actions to the set of objectively extended consequences. A state is thus given by the unique objectively extended consequence that is associated with every basic action, and the objective conceivable state space exhausts all possibilities for assigning the objectively extended consequences to the basic actions. It is uniquely and objectively defined given A and C.

Given the definition in (1), states are simply abstract notions of possible resolutions of uncertainty and do not have meaning beyond that. The state space is partitioned into a set of states $\widetilde{S} = C^A$, called fully describable, in which only known consequences may materialize, and a set of states, $S_0 \setminus \widetilde{S}$, referred to as imperfectly describable because their description involves the unknown consequence,

[4]In Karni and Vierø (2017), it was referred to as the augmented conceivable state space. The term objective conceivable state space is used in the present context to distinguish it from the subjective depiction of uncertainty below.

which cannot be fully described ex ante. Example 1 provides an illustration of the construction of the objective conceivable state space.

Example 1.a: COVID-19 vaccination decision. As a simple example of a vaccination decision, let the set of basic actions be $A = \{a_1 = \text{vaccination}, a_2 = \text{no vaccination}\}$ and the set of known consequences be $C = \{c_1 = \text{healthy}, c_2 = \text{sick}\}$. Given C, the objective unknown is defined as $x = \neg\{c_1, c_2\}$, that is, neither healthy nor sick, and the objective state space defined in (1) is given by

$$\begin{array}{cccccccccc} & s_1 & s_2 & s_3 & s_4 & s_5 & s_6 & s_7 & s_8 & s_9 \\ a_1 & c_1 & c_1 & c_1 & c_2 & c_2 & c_2 & x & x & x \\ a_2 & c_1 & c_2 & x & c_1 & c_2 & x & c_1 & c_2 & x \end{array} \tag{2}$$

The set of fully describable states is $\{s_1, s_2, s_4, s_5\}$.■

The objective conceivable state space, derived using the single objective unknown consequence x, constitutes the objective description of the world. It depicts uncertainty as it can be described and verified by an external observer.

2.2 Subjective uncertainty

The objective conceivable state space, may not reflect the problem as the decision maker sees it. The decision maker may in fact perceive a number of potential unknown consequences. They could, for example, be "something good," "something bad," something terrible," and "a minor inconvenience." These are subjective, since they incorporate how the decision maker feels about the unknown, and it cannot be verified neither what the decision maker ex ante anticipates feeling, nor how the decision maker indeed feels ex post.

Let $X_{sub} = \{x_1, \ldots, x_m\}$ be the set of subjective potential unknown consequences considered by the decision maker. Define $\widehat{C}_{sub} \equiv C \cup X_{sub}$, denoted the set of *subjectively* extended consequences. Note that the perception of consequences in X_{sub} may differ from the ex-ante and the ex-post point of view. For example, an unknown that was expected to be terrible may turn out to be a minor inconvenience once experienced. All subjectively unknown consequences in the decision maker's subjective view of the world have the same projection, to the objectively unknown consequence x, in the objective description of the world. Hence, outsiders can only verify that some unknown consequence that was not on the prior list of known consequences C materialized, but not which one of the subjective ones in X_{sub} it is.

Define the subjective conceivable state space

$$S_{sub} \equiv (\widehat{C}_{sub})^A = \{s : A \to \widehat{C}_{sub}\}. \tag{3}$$

that is, the set of all functions from the set of basic actions to the set of *subjectively* extended consequences. The subjective conceivable state space captures the perspective of the decision maker and lists the possible resolutions of uncertainty that the decision maker envisions. Since only the decision maker knows how many unknowns she perceives, only the decision maker knows which subjective state space she envisions. Like the objective conceivable state space, the subjective conceivable state space is partitioned into fully describable states in which only known consequences may materialise, and states in which new, previously unknown, consequences may be discovered and materialise. The set of fully describable states is the same as for the objective conceivable state space.

When the decision maker perceives several possible subjective unknowns, the subjective conceivable state space is a finer partition of the objective conceivable state space. Thus, the decision maker can envision more different ways that uncertainty can resolve than reflected by the objective description of uncertainty. In particular, an objective imperfectly describable state where a new consequence is discovered is partitioned into multiple subjective states, corresponding to the different perceptions of the unknown. Hence, given a particular objective state, there may still be residual subjective uncertainty, the resolution of which is not given by knowledge of the objective state. Example 1 is continued below to provide an illustration of the decision maker's subjective view of the world, captured by the subjective conceivable state space.

Example 1.b: COVID-19 vaccination decision. As before, let the set of basic actions be $A = \{a_1 = \text{vaccination}, a_2 = \text{no vaccination}\}$ and the set of known consequences be $C = \{c_1 = \text{healthy}, c_2 = \text{sick}\}$. Suppose the decision maker anticipates either a positive never heard of unknown side effect P or a negative never heard of unknown side effect N. The set of subjective unknowns is thus $X_{sub} = \{P, N\}$. Hence, there are four subjective extended consequences: $\widehat{C}_{sub} = \{c_1, c_2, P, N\}$. The subjective state space defined in (3) is given by

$$\begin{array}{ccccccccccccccccc} & s_1^s & s_2^s & s_3^s & s_4^s & s_5^s & s_6^s & s_7^s & s_8^s & s_9^s & s_{10}^s & s_{11}^s & s_{12}^s & s_{13}^s & s_{14}^s & s_{15}^s & s_{16}^s \\ a_1 & c_1 & c_1 & c_1 & c_1 & c_2 & c_2 & c_2 & c_2 & N & N & N & N & P & P & P & P \\ a_2 & c_1 & c_2 & N & P & c_1 & c_2 & N & P & c_1 & c_2 & N & P & c_1 & c_2 & N & P \end{array} \tag{4}$$

As state spaces (2) and (4) illustrate, objective states do not resolve all subjective uncertainty. While there is a unique projection of each subjective state on the objective state space, different subjective states may project to the same objective state. For example, the subjective states $s_{11}^s, s_{12}^s, s_{15}^s$ and s_{16}^s in the space (4) all project to the objective state s_9 in the space (2). Thus, knowledge of the occurrence of objective state s_9 still leaves residual uncertainty about which of the subjective states $s_{11}^s, s_{12}^s, s_{15}^s$ or s_{16}^s prevails.■

Which of the elements in X_{sub} obtains is the decision maker's private information. Therefore, strict subsets of X_{sub} cannot be contracted or bet on. As a result, nor can individual subjective imperfectly describable states. The objective unknown x and objective states do not suffer from these problems.

2.3 Objects of choice

In order to apply the revealed preference approach, we must be able to both describe bets, trades, and contracts ex ante and settle them once uncertainty resolves. Therefore, the external observer's model of the decision maker, as well as the axiomatic choice-tests, must be phrased in objective language. At the same time, we need to allow for the decision maker's subjective perception of multiple unknowns to have an effect on behavior.

The set of basic actions A naturally form part of the choice set. The elements of A can be objectively described, since they are just names of possible courses of action. For example, an element of A could be "get a Pfizer vaccine," which is well-defined and verifiable. At the same time, an element of A may include subjective elements. For example, "experience a terrible, previously unseen or unheard of side effect," and "experience a wonderful, previously unseen or unheard of side effect."

In addition to choosing basic actions, it is assumed that the decision maker can make bets on objective conceivable states, with payoffs in known consequences. The set of such bets is given by

$$B_C \equiv \{b : S_0 \to C\}.$$

These bets are state-contingent and are fully objective. Including the bets among the objects of choice serves a purpose similar to that of the objective probabilities in Anscombe and Aumann, since the bets are used to calibrate the decision maker's beliefs.

The DM is also assumed to be able to randomize among and between basic actions and state-contingent bets. Denote by $\mathscr{H} \equiv \Delta(A \cup B_C)$ the set of all possible probability distributions on $A \cup B_C$. A generic element $\mu \in \mathscr{H}$ selects an act in $A \cup B_C$ according to the distribution μ, and is to be understood as an objective randomisation between basic actions and bets. Elements of $\mathscr{H}$ are referred to as mixed acts. For notational convenience, denote the support by $H \equiv A \cup B_C$. Let h denote a generic element of H. With common abuse of notation, h also denotes the mixed act that assigns probability 1 to act h. Thus, $H \subset \mathscr{H}$.

The set of mixed acts $\mathscr{H}$ is the choice set. It satisfies the requirements of being formulated using only objective language and at the same time including subjective elements. Note that $\mathscr{H}$ is a connected compact topological space.

For all $\mu, \mu' \in \mathscr{H}$ and $\alpha \in [0, 1]$, $\alpha\mu + (1 - \alpha)\mu' \in \mathscr{H}$ is a point-wise mixture on the support of the mixed acts. It can be interpreted as a coin-flip determining which randomisation device to use. With this mixture operation, $\mathscr{H}$ is a convex set. Note that $(\alpha\mu + (1 - \alpha)\mu')(h) = \alpha\mu(h) + (1 - \alpha)\mu'(h)$ for all $h \in H$.

The decision maker has a preference relation $\succsim$ on the set of mixed acts $\mathscr{H}$. Define $\sim$ and $\succ$ as, respectively, the symmetric and asymmetric parts of $\succsim$.

2.4 Anscombe-Aumann acts

It is standard in modern decision theory to model the choice set as a set of Anscombe-Aumann (1963) acts. To obtain a model that is close in formulation to the standard model, mixed acts are translated to these. Karni and Vierø (2017) extended Anscombe-Aumann acts to potentially return a previously unknown consequence, with the consistency restriction that if a state does not prescribe a surprise for any basic action, then the Anscombe-Aumann acts cannot yield a surprise in that state either.[5] Formally, the set of objective extended Anscombe-Aumann acts is given by

$$\mathcal{F}_0 \equiv \{f : S_0 \to \Delta(\widehat{C}) \text{ such that } f(s) \in \Delta(C)\ \forall s \in \widetilde{S}\}.$$

There is a unique translation of mixed acts in $\mathscr{H}$ to objective extended Anscombe-Aumann acts. Given the definition of the objective state space, each basic action $a \in A$ maps uniquely to a state-contingent act, namely a^* such that for all $s \in S_0$, $a^*(s) = s(a) \in \widehat{C}$. In matrix (2), for example, act a_1 maps to a^* for which $a^*(s_1) = a^*(s_2) = a^*(s_3) = c_1$, $a^*(s_4) = a^*(s_5) = a^*(s_6) = c_2$, and $a^*(s_7) = a^*(s_8) = a^*(s_9) = x$. The bets specify state-contingent outcomes by definition. Thus, all acts in the support of the mixed acts have a unique translation to objective state-contingent acts that assign some consequence to each state.

For general mixed acts $\mu \in \mathscr{H}$, the translation mapping $T : \mathscr{H} \to \mathcal{F}_0$ in Definition 1 below translates each mixed act μ in $\mathscr{H}$ to a unique objective extended Anscombe-Aumann act, which returns a particular lottery $T_s(\mu)$ over objectively extended consequences if nature chooses state $s \in S_0$.

Definition 1. *The translation mapping $T : \mathscr{H} \to \mathcal{F}_0$ is defined by for all $s \in S_0$, for all $\hat{c} \in \widehat{C}$, for all $\mu \in \mathscr{H}$,*

$$T_s(\mu)(\hat{c}) = \sum_{h \in supp(\mu) : h(s) = \hat{c}} \mu(h). \tag{5}$$

Note that the translation mapping is linear, that is,

$$T(\alpha\mu + (1-\alpha)\mu') = \alpha T(\mu) + (1-\alpha)T(\mu'). \tag{6}$$

It is important to note that some aspects that form part of the decision maker's perception of the mixed acts are lost in translation, because the translation suppresses the subjective aspects. Different acts in $\mathscr{H}$ that involve different perceptions of the unknown translate to the same objective extended Anscombe-Aumannn act. That is, T^{-1} is a correspondence. However, different acts in $T^{-1}(f)$ need not be indifferent to each other.

In summary, an outside observer who seeks a tractable state-contingent model and is constrained by the need to frame the model in objective language faces a

[5] The consistency restriction ensures adherence to the revealed preference methodology.

mismatch between the objects of choice as the decision maker sees them and forms preferences over, and the objects as the outsider observes them. The mismatch provides a challenge to the revealed preferences approach. It raises the question whether we can model a decision maker who has preferences over acts that have subjective unknown elements with a state-contingent model framed in objective language. More precisely, can we model a decision maker who has preferences over the mixed acts in $\mathscr{H}$, as if she is maximising expected utility over the translations of these to objective extended Anscombe-Aumann acts, without surpressing that the subjective elements impact behavior?

3 Seemingly irrational behaviour

As discussed, outside observers need to rely on a model that is framed in objective language, because only resolutions of uncertainty that can be objectively described are verifiable and tradable. However, if we fail to recognize that a decision maker perceives multiple subjective unknowns, it poses a serious problem. A rational decision maker may display seemingly irrational behaviour. Suppose that a decision maker considers multiple subjective unknowns and views the acts in her choice set to potentially result in different of these. Choices will appear inconsistent and thus irrational in the reduced objective world given by the objective state space and objective state-contingent acts, when they are in fact rational given the decision maker's extended language and view of the problem. The following example illustrates the issue in the context of a vaccination decision.

Example 2: Suppose we are interested in eliciting a decision maker's preferences over getting a COVID-19 vaccine vs. not getting a vaccine. Two different COVID-19 vaccines are available, which have identical scientific descriptions and only differ in their names. However, the decision maker perceives one vaccine to potentially result in a positive, never seen or heard of, unknown consequence and the other in a negative, also never seen or heard of, unknown consequence, and as a result has the following rankings:

$$\text{Vaccine 1} \succ \text{No vaccine}$$

and

$$\text{No vaccine} \succ \text{Vaccine 2}$$

Suppose the outside observer does not realise that the decision maker perceives the vaccines as resulting in different unknowns and therefore simply considers the identical objective descriptions, and as a result categorizes each of the two vaccines simply as "vaccine". Then the choice between "vaccine" and "no vaccine" will appear to violate rationality, because it violates either transitivity or asymmetry of the strict preference relation:

$$\text{Vaccine} \succ \text{No vaccine} \succ \text{Vaccine}$$

However, the violation is due to deficiencies in properly modelling the problem rather than irrationality of the decision maker.

The example can be elaborated upon by considering the state spaces in Example 1. Reinterpret the basic actions as a_1=Vaccine 1 and a_2=Vaccine 2, with $T_{s_1}(a_1) = T_{s_1}(a_2)(s_1) = c_1$, $T_{s_5}(a_1) = T_{s_5}(a_2) = c_2$, and $T_{s_9}(a_1) = T_{s_9}(a_2) = x$. Assume that beliefs over the objective state space (2) are given by the distribution π, with $\pi(s_2) = \pi(s_3) = \pi(s_4) = \pi(s_6) = \pi(s_7) = \pi(s_8) = 0$.

Vaccine 1 and Vaccine 2 thus have the same objective translation. With probability $\pi(s_1)$, they result in c_1, with probability $\pi(s_5)$, they result in c_2, and with probability $\pi(s_9)$, they result in something new. Thus, they look identical to the analyst. Call this objective translation "Vaccine". The axioms in Karni and Vierø (2017) would imply that $a_1 \sim a_2$.

Now suppose the decision maker perceives of two subjective unknowns, P and N, with $P \succ N$. In state spaces (2) and (4), $s_1 = s^s_1$, $s_5 = s^s_6$, and $s_9 = \{s^s_{11}, s^s_{12}, s^s_{15}, s^s_{16}\}$. If the decision maker bases his decisions on the subjective view in (4), and he perceives s^s_{15} more likely than s^s_{12}, then $a_1 \succ a_2$. But if the analyst works with the objective translations and the axiomatization in Karni and Vierø (2017), it looks like Vaccine$\succ$Vaccine, such that choices appear to violate rationality. ■

As Example 2 demonstrates, behaviour that is rational given the extended language of the decision maker may look irrational when described with the reduced language of an analyst, who describes the objects of choice as objective state-contingent acts returning lotteries with known consequences and a single unknown, if the analyst fails to recognize that subjective refinements of the unknown affect behavior. A good model therefore needs to accommodate that subjective perceptions affect behavior, but that the subjective features are lost in translation to the objective description.[6]

4 Preference structure

The main representation result provides conditions under which we can use a model with objective extended Anscombe-Aumann acts and an expected utility representation over these to capture the behaviour of a decision maker who forms preferences over mixed acts that have subjective unknown elements. As in Karni and Vierø (2017), the expected utility representation includes the decision maker's attitude towards the unknown. However, in order to avoid the seeming violations of rationality due to losses in translation from subjective to objective acts, the decision maker's attitude towards the unknown is both state-dependent and act-dependent.

The preference relation is assumed to satisfy the following standard axioms:

Axiom 1 (Pre-order)**.** *The strict preference relation $\succ$ on $\mathscr{H}$ is asymmetric and negatively transitive.*

Axiom 2 (Continuity)**.** *For all $\mu \in \mathscr{H}$, the sets $\{\mu' | \mu' \succsim \mu\}$ and $\{\mu' | \mu \succsim \mu'\}$ are closed.*

Axiom 3 (Independence)**.** *For all $\mu, \mu', \mu'' \in \mathscr{H}$ and $\alpha \in (0, 1]$, $\mu \succ \mu''$ implies $\alpha\mu + (1-\alpha)\mu' \succ \alpha\mu'' + (1-\alpha)\mu'$.*

4.1 Conditional preferences, null states, and notations

It is assumed that the decision maker is able to express ex-ante conditional preferences. These express the decision maker's opinion about which one of two acts in $\mathscr{H}$ will lead to more desirable residual uncertainty if any given objective event $E \subseteq S_0$ occurs. The residual uncertainty may be about objective states if E is non-degenerate, or about subjective states if E contains imperfectly describable states. It is important to emphasize the ex-ante perspective, since the actual resolution of uncertainty may reveal new consequences, and the decision maker then

[6]Further difficulty would arise at the interpersonal level: a particular objective state-contingent act may be a translation of different subjective acts for different decision makers.

experiences how the new consequences actually affect her well-being. Hence, the perception of the consequences and thus the ranking of acts may change ex-post, such that ex-post preferences may not agree with ex-ante conditional preferences.

For any $E \subseteq S_0$ and for any $\mu, \lambda \in \mathscr{H}$, the conditional mixed act $\mu_E\lambda$ returns the outcome of mixed act μ in event E and the outcome of mixed act λ in event $S_0 \setminus E$. Fix a mixed act $\lambda \in \Delta(B_C)$.

Definition 2 (Ex-ante conditional preferences). *For any event $E \subseteq S_0$, ex-ante preferences conditional on E are given by $\mu \succsim_E \mu'$ if $\mu_E\lambda \succsim \mu'_E\lambda$.*

An objective state s is said to be null if $\mu \sim_s \mu'$ for all $\mu, \mu' \in \mathscr{H}$. An objective state is said to be nonnull if it is not null.

Define the following notations: Let

$$\mathcal{F}_C = \{f : S_0 \to \Delta(C)\}$$

be the set of objective Anscombe-Aumann acts that return lotteries over *known* consequences, i.e. lotteries with support a subset of C. Notice that

$$\Delta(B_C) = \{\mu \in \mathscr{H} | T(\mu) \in \mathcal{F}_C\}.$$

Thus, the set of mixed acts whose support is a subset of B_C is exactly the set of mixed acts whose translations return lotteries over known consequences only. Notice also that the translation mapping from $\Delta(B_C)$ to $\mathcal{F}_C$ is onto.

Let

$$L = \{\mu \in \mathscr{H} | T_s(\mu) = T_{s'}(\mu) \in \Delta(C)\ \forall s, s' \in S_0\}$$

be the set of mixed acts that translate to constant Anscombe-Aumann acts in $\mathcal{F}_C$. Note that $L \subset \Delta(B_C)$ and $\Delta(B_C) \subset \mathscr{H}$.

Let $l_p \in L$ denote an act that translates to the constant Anscombe-Aumann act $p \in \Delta(C)$ and let L_p denote the set of these acts, that is, the set of all acts that translate to the same constant Anscombe-Aumann act p.

4.2 Further preference structure and main representation result

The next few axioms are only imposed on preferences over $\Delta(B_C)$, that is, on preferences over mixed acts that do not reveal new consequences under any circumstance. Axiom 4 states that as long as no new consequence is involved, the decision maker is indifferent between mixed acts that translate to the same objective Anscombe-Aumann act.

Axiom 4 (Extended indifference)**.** *For all $\mu, \mu' \in \Delta(B_C)$, if $T(\mu) = T(\mu')$ then $\mu \sim \mu'$.*

It is important that Axiom 4 is only imposed on $\Delta(B_C)$ rather than on preferences over the full set of mixed acts, exactly because we want to allow for the flexibility of evaluating different subjective unknowns differently. That is, Axiom 4 does not in general impose that the decision maker is indifferent between acts that translate to the same objective extended Anscombe-Aumann act. Specifically, when the objective extended Anscombe-Aumann act belongs to $\mathscr{H} \setminus \Delta(B_C)$, the decision maker is allowed to replace the objective unknown x with lotteries over subjective unknowns, with different such lotteries for different acts. Because the subjective consequences may be different for two mixed acts that translate to the same objective extended Anscombe-Aumann act, the decision maker may have a strict preference between them.

By including Axiom 4, we get separability over states.

Proposition 1. *If a preference relation $\succsim$ on $\Delta(B_C)$ satisfies Axioms 1 through 4, then there exist real-valued functions $\{W_s\}_{s \in S_0}$ on C, unique up to cardinal unit-comparable transformation, such that for all $\mu, \mu' \in \Delta(B_C)$,*

$$\mu \succ \mu' \Leftrightarrow \sum_{h \in B_C} \mu(h) \sum_{s \in S_0} W_s(T_s(h)) > \sum_{h \in B_C} \mu'(h) \sum_{s \in S_0} W_s(T_s(h)). \tag{7}$$

The proof of Proposition 1 is in the Appendix.

Like Axiom 4, the next two axioms are also only imposed on preferences over $\Delta(B_C)$.

Axiom 5 (Conditional non-triviality)**.** *The strict conditional preference relation $\succ_s$ on $\Delta(B_C)$ is non-empty for at least three states s.*

Note that when $\succ_s$ is non-empty on $\Delta(B_C)$, it implies that $\succ_s$ is nonempty on $\mathscr{H}$.

Axiom 6 ($\Delta(B_C)$-Monotonicity)**.** *For all $p, q \in \Delta(C)$, for all non-null states s, for all $\mu, \mu' \in \Delta(B_C)$ with $T_{s'}(\mu) = T_{s'}(\mu')$ for all $s' \neq s$, $T_s(\mu) = p$, $T_s(\mu') = q$, we have $\mu \succsim \mu' \Leftrightarrow l_p \succsim l_q$ for all $l_p \in L_p$ and $l_q \in L_q$.*

Axiom 6 states that if two mixed acts in $\Delta(B_C)$ only differ in the residual uncertainty in objective state s, and that residual uncertainty translates to lotteries p and q, then the ranking of these two acts is the same as the ranking of acts that translate to those lotteries everywhere.

By adding Axioms 5 and 6, we can obtain subjective probabilities of objective states in S_0 and state-independent Bernoulli utilities of known consequences (see Lemma 3 in the Appendix). We thus have a standard subjective expected utility representation over objective extended Anscombe-Aumann acts when we restrict attention to acts that do not result in new consequences. This provides structure for the remaining task, since it provides a measuring rod against which we can compare the anticipated well-being from unknown consequences.

The remaining task is to obtain a utility representation over all acts, which has an expected utility structure but still allows the decision maker to treat different subjective unknowns differently. The following axioms are imposed on preferences over all acts, including those that may return a yet unknown consequence.

Axiom 7 (Coherence)**.** *For any disjoint events $E, G \subseteq S_0$ and mixed acts $\mu, \mu' \in \mathscr{H}$,*

1. *$\mu \succsim_E \mu'$ and $\mu \succsim_G \mu'$ implies $\mu \succsim_{E \cup G} \mu'$*
2. *$\mu \succ_E \mu'$ and $\mu \succsim_G \mu'$ implies $\mu \succ_{E \cup G} \mu'$.*

Axiom 7 postulates that if the residual uncertainty of μ is preferred over that of μ' in both events E and G, then it is also the case in the combined event. Note that given Axiom 7, Axiom 5 also implies non-triviality of the unconditional preference relation $\succ$.

Axiom 8 (Separability)**.** *For all events $E \subseteq S_0$ and for all $\mu, \mu' \in \mathscr{H}$ for which $\mu \sim_{S_0 \setminus E} \mu'$, it holds that $\mu \succsim \mu' \Leftrightarrow \mu \succsim_E \mu'$.*

Axiom 8 concerns acts for which the residual uncertainty in event $S_0 \setminus E$ is considered to be equally good from an ex-ante point of view and says that the decision maker's ranking of such acts is given by her ex-ante ranking of the residual uncertainty of the acts in event E.

Theorem 1 below is the main representation result. It shows that we can model a decision maker who perceives multiple subjective unknowns as a generalized expected utility maximizer, with a Bernouilli utility function over known consequences, subjective probabilities over objective states, and a state- and act-dependent attitude towards the unknown.

Theorem 1. *If $\succ$ satisfies Axioms 1 through 8, there exist functions $u : C \to \mathbb{R}$ and $u_x : (S_0 \setminus \widetilde{S}) \times \mathscr{H} \to \mathbb{R}$ and a probability measure π on S_0 such that for all $\mu, \mu' \in \mathscr{H}$,*

$$\sum_{s\in S_0} \pi(s) \left[\sum_{c\in C} T_s(\mu)(c)u(c) + \Big(1 - \sum_{c\in C} T_s(\mu)(c)\Big) u_x(s,\mu) \right]$$
$$\geq \sum_{s\in S_0} \pi(s) \left[\sum_{c\in C} T_s(\mu')(c)u(c) + \Big(1 - \sum_{c\in C} T_s(\mu')(c)\Big) u_x(s,\mu') \right]. \tag{8}$$

The function u is unique up to positive linear transformations, and the function u_x is unique given the transformation of u. Moreover, the probability measure π is unique, and $\pi(s) = 0$ if and only if s is null.

The proof of Theorem 1 is in the Appendix.

For each objective state, a mixed act μ is translated to a lottery $T_s(\mu)$ over $\widehat{C}$. The representation evaluates that lottery by computing a generalized von Neumann-Morgenstern utility, which evaluates all known outcomes according to the Bernoulli utility function $u(\cdot)$, and evaluates the unknown objective outcome x according to the act- and state dependent function $u_x(s,\mu)$. The translation assigns probability mass $1 - \sum_{c\in C} T_s(\mu)(c)$, that is, the residual probability not assigned to known outcomes, to the unknown outcome.

The function $u_x(s,\mu)$ captures how the decision maker feels about unknowns potentially arising from act μ in objective state s. It is the decision maker's act- and state dependent attitude towards the unknown. The state dependence arises because the decision maker can think of act μ as resulting in different subjective unknowns in different states. If $u_x(s,\mu) < u_x(s',\mu)$, then act μ is perceived as resulting in a more adverse discovery in state s than in state s'. The act dependence arises because the decision maker can think of different actions resulting in different subjective unknowns. If $u_x(s,\mu) < u_x(s,\mu')$, then act μ is perceived as resulting in a more adverse discovery than μ' in state s. If the decision maker perceives a single unknown, then the function $u_x(s,\mu)$ is constant. In that case, acts are evaluated as in Karni and Vierø (2017).

Since the proof of Theorem 1 applies a theorem from Skiadas (1997a), it is worth mentioning how the present axioms relate to his. Some of Skiadas' axioms are implied by the present axioms. Other of Skiadas' axioms are redundant because the present set of acts is a mixture space. Also, because the residual uncertainty given an objective state is only about the unknown consequences, a lot of structure is obtained from the axioms imposed on preference over $\Delta(B_C)$. Axiom 7 is also imposed by Skiadas.

5 How many unknowns?

Section 3 established the importance of acknowledging that the decision maker potentially perceives multiple subjective unknowns. It is therefore important to establish how one can elicit whether a decision maker indeed has such a perception. We can directly use that asymmetry of the strict preference relation will be violated if we model the decision maker with a single unknown, when he in fact perceives multiple unknowns. The following asymmetry test can be used to reveal the decision maker's perception:

Definition 3 (Asymmetry test to reveal multiple subjective unknowns:)**.** *Does there exist $s \in S_0$ and $\mu, \lambda \in \mathscr{H}$, with $T_{s'}(\mu) = T_{s'}(\lambda) \in C$ for $s' \neq s$ and $T_s(\mu) = T_s(\lambda) = x$ such that $\mu \succ \lambda$?*

If the answer to the question in Definition 3 is yes, we need to model the decision maker with multiple subjective unknowns.

In Definition 3, the mixed acts μ and λ translate to the same known consequence in all states except from s. In state s, both acts translate to the objective unknown. If the decision maker prefers one act over the other, it must be because the unknown arising from μ is perceived by the decision maker to be different from the unknown arising from λ.

A positive insight is that the problem with seemingly irrational choice can be addressed by including multiple subjective unknowns. Apparent violations of rationality, that are due to deficiencies in properly modelling the problem, disappear when the problem is modelled properly.

Example 3: Can we distinguish how many multiples? The decision maker in general considers a non-degenerate set of subjective unknowns among the possible resolutions of uncertainty. The corresponding subjective state space is S_{sub}, defined in (3). An objective state $s_o \in S_0 \setminus \tilde{S}$ is the projection of a non-degenerate event $E(s_o)$ in S_{sub}. For example, in state spaces (2) and (4), $E(s_7) = \{s_9^s, s_{13}^s\}$. The objective consequence x in state s_o is the translation of (mixtures of) subjective consequences in the subjective states in $E(s_o)$. In other words, if $f(s_o) = x$, then $T^{-1}(f)$ may involve any combination of the subjective unknowns in the states in $E(s_o)$.

In the decision maker's view of the world, the proper translation of mixed acts would be to *subjective* extended Anscombe-Aumann acts, given by

$$\mathcal{F}_{sub} \equiv \{f : S_{sub} \to \Delta(\widehat{C}_{sub}) \text{ such that } f(s) \in \Delta(C)\ \forall s \in \widetilde{S}\}.$$

They assign a lottery $q_{s'}$ with support a subset of $\widehat{C}_{sub}$ to each subjective state $s' \in E(s_o)$.

Suppose the decision maker acts as if he were an expected utility maximizer over subjective states and translations T_{sub} of acts to subjective extended Anscombe-Aumann acts, with subjective probabilities $\hat{\pi}$ and utilities $\hat{u} : \widehat{C}_{sub} \to \mathbb{R}$. Then we would have that the attitude toward the unknown in (8) is

$$u_x(s_o, \mu) = \sum_{s' \in E(s_o)} \hat{\pi}_{s'} \sum_{x_i \in X_{sub}} T_{sub}(s')(x_i)\hat{u}(x_i).$$

It will now be illustrated that while we can elicit whether the decision maker perceives one or multiple unknowns using the test in Definition 3, we cannot distinguish how many multiples. To see this, assume that $A = \{a_1, a_2\}$ and $C = \{c\}$.

Then the objective conceivable state space is $\widehat{C}^A$, given by

$$\begin{array}{ccccc} & s_1 & s_2 & s_3 & s_4 \\ a_1 & c & c & x & x \\ a_2 & c & x & c & x \end{array}$$

Suppose that the decision maker perceives three subjective unknowns $x_3 \succ x_2 \succ x_1$. In the corresponding subjective state space $(\widehat{C}_{sub})^A$, consider the event

$$E(s_2) = \left\{ \begin{pmatrix} c \\ x_1 \end{pmatrix}, \begin{pmatrix} c \\ x_2 \end{pmatrix}, \begin{pmatrix} c \\ x_3 \end{pmatrix} \right\}.$$

Assume the decision maker translates actions to subjective state contingent acts that return lotteries over known consequences and the three subjective unknowns and that he acts as if he is an expected utility maximiser. For an act μ, let $q_s(x_i) = T_{sub}(\mu)(s)(x_i)$. Then the attitude toward the unknown that we observe in the objective world is $u_x(s_2, \mu) = U_{E(s_2)} = \sum_{s=1}^{3} \hat{\pi}_s \sum_{i=1}^{3} q_s(x_i)\hat{u}(x_i)$. Since Archimedean continuity is implied by Axiom 2, any intermediate unknown is indifferent to a mixture of the best and worst unknowns. I.e. there exists $\alpha \in [0, 1]$ such that $x_2 \sim \alpha x_3 + (1 - \alpha)x_1$. Hence,

$$U_{E(s_2)} = \sum_{s=1}^{3} \hat{\pi}_s[(q_s(x_1) + q_s(x_2)(1 - \alpha))\hat{u}(x_1) + (q_s(x_3) + q_s(x_2)\alpha)\hat{u}(x_3)].$$

Consider instead a decision maker who only perceives of the two subjective unknowns, $x_3 \succ x_1$, but is otherwise identical to the previous decision maker. In the corresponding subjective state space, consider the event

$$E'(s_2) = \left\{ \begin{pmatrix} c \\ x_1 \end{pmatrix}, \begin{pmatrix} c \\ x_3 \end{pmatrix} \right\},$$

and let the subjective probabilities of the two states in $E'(s_2)$ be π'_1 and π'_3, respectively. For an act μ, let $q'_s(x_i) = T_{sub}(\mu)(s)(x_i)$ for this decision maker. Then

$$u_x(s_2, \mu) = U_{E'(s_2)} = \pi'_1\big(q'_1(x_1)\hat{u}(x_1) + q'_1(x_3)\hat{u}(x_3)\big) + \pi'_3\big(q'_3(x_1)\hat{u}(x_1) + q'_3(x_3)\hat{u}(x_3)\big).$$

We see that $U_{E(s_2)} = U_{E'(s_2)}$ for all $\hat{u}$ if $\pi'_1 q'_1(x_1) + \pi'_3 q'_3(x_1) = \sum_{s=1}^{3} \hat{\pi}_s[q_s(x_1) + q_s(x_2)(1 - \alpha)]$. Hence the two decision makers are indistinguishable.■

By induction, a similar argument to that in Example 3 can be repeated for any finite number of subjective unknowns. Thus, due to the lack of structure of not observing the underlying set of unknowns and corresponding subjective state-contingent acts, the decision analyst cannot distinguish behaviour when the decision maker considers $m > 1$ subjective unknowns from behaviour when he considers $m' > 1$ subjective unknowns, where $m \neq m'$. The set of utilities of subjective unknowns is spanned by the utilities of the very best and very worst subjective unknowns. Thus, the decision analyst can reveal that there are multiple unknowns using the asymmetry test in Definition 3, but not how many. In principle, since u_x is act dependent, there may be as many subjective unknown as there are acts.

The inability of the analyst to distinguish how many unknowns the decision maker perceives comes from the underlying subjectivity of both states and acts. When the analyst models the decision maker using the objective state space and objective extended Anscombe-Aumann acts defined on that, the analyst models both the resolutions of uncertainty and the objects of choice coarser than the decision maker considers them. The issue also stems from convexifying the choice set. The upside is that it is sufficient to use two subjective unknowns to model the decision maker: The very best and very worst span the others. Alternatively to the usual coin-flip interpretation, a convex combination of the very best and very worst unknowns is a proxy for the perceived attractiveness of the unknown.

6 Discussion

This paper has argued for the importance of considering the potential for multiple subjective unknowns when the decision maker is aware of her own unawareness. It has shown that new issues arise due to the subjectivity, which complicates modelling-steps that are usually taken without much consideration. Part of the decision maker's perception of the problem that influences her preferences is lost in translation in the objective modelling of the problem.

In standard modelling of decisions under uncertainty, the events relevant to both the decision maker and the analysis are modelled with a state space. Consequently, the analyst is aware of everything the decision maker is aware of and vice versa. This is not the case when unawareness is studied.

With unawareness, there are different perspectives on the problem. First, there is the decision maker's perspective, which is given by a subjective state-space. Second, there is the perspective of the decision theorist, who can take an abstract meta-view and envision two state-spaces, the richer subjective space of the decision maker, and a poorer objective space. The third perspective is that of the decision analyst, or outside observer, who is interested in learning from choices within the objective perspective in order to describe and predict behavior. In traditional decision theory the perspectives of the decision theorist and the decision analyst are the same. In the present context of subjective unknowns, they differ. The decision analyst lacks the full perspective of the decision theorist.

Thus, we are considering a decision maker, who makes choices based on her subjective view of the world. We also have a decision analyst, who seeks to learn from the choices of the decision maker. The present paper has taken the meta-view of the decision theorist and provided an answer to the question of what the decision analyst can learn from the decision maker's choices.

The paper takes choice over courses of action as the primitive. The courses of action are translated to objective Anscombe-Aumann acts. The axiomatic structure implies the existence of an expected utility representation that contains an act- and state-dependent attitude towards the unknown that captures the decision maker's subjective perception of discoveries that may result from choosing a particular course of action. The results enable tractable modelling of decision makers, while at the same time allowing that different subjective perceptions of the unknown in different situations affect behaviour.

Appendix

Lemma 1. *If the preference relation $\succ$ satisfies Axioms 1 through 3, then there exists a function $\mathcal{V} : \mathscr{H} \to \mathbb{R}$ such that for all $\mu, \mu' \in \mathscr{H}$, $\mu \succ \mu' \Leftrightarrow \mathcal{V}(\mu) > \mathcal{V}(\mu')$. Furthermore, $\mathcal{V}(\alpha\mu + (1-\alpha)\mu') = \alpha\mathcal{V}(\mu) + (1-\alpha)\mathcal{V}(\mu')$, and $\mathcal{V}$ is unique up to positive affine transformation.*

Proof: The set of mixed acts $\mathscr{H}$ is a convex subset of a Euclidean space. Therefore, by Proposition 2 of Uyanik and Khan (2022), Axiom 2 (Continuity) implies the Archimedean axiom.[7] Then the mixture space theorem immediately implies existence of a representation as well as linearity.

Lemma 2. *If the preference relation $\succ$ satisfies Axioms 1 through 3, then there exists a function $V : H \to \mathbb{R}$ such that for all $\mu, \mu' \in \mathscr{H}$,*

$$\mu \succ \mu' \Leftrightarrow \sum_{h\in H} \mu(h)V(h) > \sum_{h\in H} \mu'(h)V(h). \tag{9}$$

Moreover, V is unique up to positive affine transformation.

Since the proof of Lemma 2 follows standard arguments, it is omitted.

Proof of Proposition 1 Let $n = |S_0|$ denote the cardinality of S_0. Fix $h^* \in \Delta(B_C)$. For each $h \in \Delta(B_C)$ and $s \in S_0$, let $f^s_h \in \mathcal{F}_C$ be defined by $f^s_h(s) = T_s(h)$ and $f^s_h(s') = T_{s'}(h^*)$ if $s' \neq s$. Let $f^* = T(h^*)$.

Consider the mixed act $\mu \in \Delta(B_C)$ that assigns probability $\frac{1}{n}$ to h and probability $\frac{n-1}{n}$ to h^*. Since the translation mapping from $\Delta(B_C)$ to $\mathcal{F}_C$ is onto, there exists a mixed act $\mu' \in \Delta(B_C)$ with translation $T(\mu')(s') = \frac{1}{n}\sum_{s\in S_0} f^s_h(s')$.

By property (6), linearity,

$$T(\mu) = T(\mu').$$

Thus, by Axiom 4, $\mu \sim \mu'$. By the representation in Lemma 1, the indifference $\mu \sim \mu'$ is equivalent to

$$\frac{1}{n}\mathcal{V}(h) + \frac{n-1}{n}\mathcal{V}(h^*) = \mathcal{V}(\mu') \tag{10}$$

For $p \in \Delta(C)$ and $f \in \mathcal{F}_C$, let $p_s f \in \mathcal{F}_C$ denote the Anscombe-Aumann act that returns p in state s and agrees with f in all $s' \neq s$.

For each $s \in S_0$, define $W_s(\cdot) : \Delta(C) \to \mathbb{R}$ by

$$W_s(p) = \mathcal{V}(T^{-1}(p_s f^*)) - \frac{n-1}{n}\mathcal{V}(h^*) \tag{11}$$

The mixed act $T^{-1}(p_s f^*)$ exists since the translation mapping from $\Delta(B_C)$ to $\mathcal{F}_C$ is onto. Thus, for $h \in H_C$,

$$W_s(T_s(h)) = \mathcal{V}(T^{-1}(T_s(h)_s f^*)) - \frac{n-1}{n}\mathcal{V}(h^*) = \mathcal{V}(T^{-1}(f^s_h)) - \frac{n-1}{n}\mathcal{V}(h^*).$$

[7] The stronger form of Continuity is needed for later results.

This implies that

$$\sum_{s\in S_0} W_s(T_s(h)) = \sum_{s\in S_0} \mathcal{V}(T^{-1}(f_h^s)) - (n-1)\mathcal{V}(h^*).$$

Multiplying by $\frac{1}{n}$ on both sides yields

$$\frac{1}{n}\sum_{s\in S_0} W_s(T_s(h)) = \frac{1}{n}\sum_{s\in S_0} \mathcal{V}(T^{-1}(f_h^s)) - \frac{n-1}{n}\mathcal{V}(h^*). \tag{12}$$

Notice that $\frac{1}{n}\sum_{s\in S_0} \mathcal{V}(T^{-1}(f_h^s)) = \mathcal{V}(\frac{1}{n}\sum_{s\in S_0} T^{-1}(f_h^s))$ by Lemma 1, and that

$$T\Big(\frac{1}{n}\sum_{s\in S_0} T^{-1}(f_h^s)\Big) = \frac{1}{n}\sum_{s\in S_0} f_h^s = T(\mu')$$

by property (6). Thus, applying Axiom 4 again and combining (10) and (12) gives $\mathcal{V}(h) = \sum_{s\in S_0} W_s(T_s(h))$. Define $V(h) \equiv \mathcal{V}(h)$. Now, plugging into (9) in Lemma 2 gives the result.

The uniqueness of $\{W_s\}_{s\in S_0}$ follows from that of V. To see this, define, for all $s \in S_0$,

$$\hat{W}_s(\cdot) = bW_s(\cdot) + d_s,\ b > 0.$$

By definition in (11), we have that for all $s \in S_0$ and $p \in \Delta(C)$,

$$\hat{W}_s(p) = b\big[V(T^{-1}(p_s f^*)) - \frac{n-1}{n}V(h^*)\big] + d_s.$$

It follows that

$$\sum_{s\in S_0} \hat{W}_s(T(h)(s)) = b\sum_{s\in S_0} V(h^s) - (n-1)V(h^*) + d = bV(h) + d,$$

where $d = \sum_{s\in S_0} d_s$. Since V is unique up to positive linear transformation, $\hat{V} = bV + d$ represents the same preferences as V. It follows that $\{\hat{W}_s\}_{s\in S_0}$ represents the same preferences as $\{W_s\}_{s\in S_0}$. ■

Lemma for main result

Lemma 3. *If the preference relation on $\Delta(B_C)$ satisfies Axioms 1 through 6, then there exists a real-valued, continuous, non-constant, affine function U on $\Delta(C)$ and a probability measure π on S_0, such that for all $\mu, \mu' \in \Delta(B_C)$,*

$$\mu \succ \mu' \Leftrightarrow \sum_{s\in S_0} \pi(s)U(T_s(\mu)) > \sum_{s\in S_0} \pi(s)U(T_s(\mu')). \tag{13}$$

Moreover, the function U is unique up to positive linear transformation, the probability measure π is unique, and $\pi(s) = 0$ if and only if s is null.

Proof of Lemma 3 By Proposition 1,

$$\mu \succ \mu' \Leftrightarrow \sum_{h\in B_C} \mu(h) \sum_{s\in S_0} W_s(T_s(h)) > \sum_{h\in B_C} \mu'(h) \sum_{s\in S_0} W_s(T_s(h)). \tag{14}$$

Fix an objective Anscombe-Aumann act $f^* \in \mathcal{F}_C$. For $p, q \in \Delta(C)$, let $\mu_\alpha \in \Delta(B_C)$ be such that $T(\mu_\alpha) = (\alpha p + (1-\alpha)q)_s f^*$, let $\mu_p \in \Delta(B_C)$ be such that $T(\mu_p) = p_s f^*$, and let $\mu_q \in \Delta(B_C)$ be such that $T(\mu_q) = q_s f^*$. These mixed acts exist since the translation mapping from $\Delta(B_C)$ to $\mathcal{F}_C$ is onto. Consider the mixed act $\alpha\mu_p + (1-\alpha)\mu_q$. By property (6), linearity of the translation mapping, $T(\alpha\mu_p + (1-\alpha)\mu_q) = \alpha T(\mu_p) + (1-\alpha)T(\mu_q) = (\alpha p + (1-\alpha)q)_s f^* = T(\mu_\alpha)$. Thus, by Axiom 4,

$$\mathcal{V}(\mu_\alpha) = \alpha\mathcal{V}(\mu_p) + (1-\alpha)\mathcal{V}(\mu_q). \tag{15}$$

We can therefore define $W_s(c) \equiv W_s(\delta_c)$ and use the usual induction argument to show that $W_s(T_s(h)) = \sum_{c\in C} T_s(h)(c)W_s(c)$. Thus,

$$\sum_{h\in B_C} \mu(h) \sum_{s\in S_0} W_s(T_s(h)) = \sum_{h\in B_C} \mu'(h) \sum_{s\in S_0}\sum_{c\in C} T_s(h)(c)W_s(c). \tag{16}$$

Consider $h, h' \in B_C$ and recall that $B_C \subset \Delta B_C$. Let $\mu_\beta = \beta h + (1-\beta)h'$. By (16),

$$\begin{aligned}\mathcal{V}(\mu_\beta) &= \beta \sum_{s\in S_0}\sum_{c\in C} T_s(h)(c)W_s(c) + (1-\beta) \sum_{s\in S_0}\sum_{c\in C} T_s(h')(c)W_s(c)\\ &= \sum_{s\in S_0}\sum_{c\in C} \big(\beta T_s(h)(c) + (1-\beta)T_s(h')(c)\big)W_s(c)\end{aligned} \tag{17}$$

By property (6), $T(\mu_\beta) = \beta T(h) + (1-\beta)T(h')$. Thus, the right hand side of (17) equals $\sum_{s\in S_0}\sum_{c\in C} T_s(\mu_\beta)(c)W_s(c)$.

Using the same steps in an induction argument implies that

$$\sum_{h\in B_C} \mu(h) \sum_{s\in S_0}\sum_{c\in C} T_s(h)(c)W_s(c) = \sum_{s\in S_0}\sum_{c\in C} T_s(\mu)(c)W_s(c). \tag{18}$$

Fix a non-null $s' \in S_0$, which exists by Axiom 5. Define for each $p \in \Delta(C)$, $U(p) = \sum_{c\in C} W_{s'}(c)p(c)$. By Axiom 6, for any $p, q \in \Delta(C)$,

$$\sum_{c\in C} p(c)W_{s'}(c) > \sum_{c\in C} q(c)W_{s'}(c) \Leftrightarrow \sum_{c\in C} p(c)W_s(c) > \sum_{c\in C} q(c)W_s(c) \tag{19}$$

for all non-null $s \in S_0$.

Thus, standard arguments imply that

$$\mu \succ \mu' \Leftrightarrow \sum_{s \in S_0} \pi(s)U(T_s(\mu)) > \sum_{s \in S_0} \pi(s)U(T_s(\mu')), \tag{20}$$

where $U(T_s(\mu)) = \sum_{c \in C} T_s(\mu)(c)u(c)$, where $u(c) = aW_s(c) + b$ for some nonnull $s \in S_0$.

Uniqueness of U follows from that of W_s. The argument for uniqueness of π and for $\pi(s) = 0$ if and only if s is null is standard. ■

Proof of Theorem 1 Axioms 8 and 1 through 3 imply that there exists a continuous function $\mathcal{U} : S_0 \times \mathscr{H} \to \mathbb{R}$ such that for all $\mu, \mu' \in \mathscr{H}$, for all $s \in S_0$

$$\mu \succsim_s \mu' \Leftrightarrow \mathcal{U}(\mu, s) \geq \mathcal{U}(\mu', s). \tag{21}$$

By Theorem 1 of Skiadas (1997a), Axioms 1, 2 and 7, together with the fact that $\mathscr{H}$ is a mixture space, implies that there exists an aggregator of state-conditional utilities. By the Theorem on p. 10 of Debreu (1959), Axioms 1, 2, 5, and 8, together with the fact that $\mathscr{H}$ is a mixture space, imply that the aggregator is additive. Thus, there exists a continuous and unique additive representation over $S_0 \times \mathscr{H}$, where the uniqueness is up to positive linear transformations. Therefore, we can write $\mathcal{V}(\mu) = \sum_{s \in S_0} \mathcal{U}(\mu, s)$, and we have that

$$\mu \succsim \mu' \Leftrightarrow \sum_{s \in S_0} \mathcal{U}(\mu, s) \geq \sum_{s \in S_0} \mathcal{U}(\mu', s). \tag{22}$$

By Lemma 3, when $T_s(\mu) \in \Delta(C)$, we can set $\mathcal{U}(\mu, s) = \pi(s)U(T_s(\mu))$.

Let μ_q and μ_x be as in Axiom 8 with $T_{s'}(\mu_x) = T_{s'}(\mu_q) \in \Delta(C)$ for all $s' \neq s$, $T_s(\mu_x) = x$, and $T_s(\mu_q) = q \in \Delta(C)$. By Proposition 1, $\mu \succ \mu' \Leftrightarrow \mathcal{V}(\mu) > \mathcal{V}(\mu')$ for any mixed acts μ, μ', and by Lemma 3, we can write $\mathcal{V}(\mu_q) = \sum_{s \in S_0} \pi(s)U(T_s(\mu_q))$.

By Proposition 1,

$$\begin{aligned}
\mathcal{V}(\alpha\mu_q + (1-\alpha)\mu_x) &= \alpha\mathcal{V}(\mu_q) + (1-\alpha)\mathcal{V}(\mu_x) \\
&= \alpha \sum_{s' \in S_0} \pi(s')U(T_{s'}(\mu_q)) + (1-\alpha) \sum_{s' \in S_0} \mathcal{U}(\mu_x, s') \\
&= \sum_{s' \neq s} \pi(s')U(T_{s'}(\mu_q)) + \alpha\pi(s)U(T_s(\mu_q)) + (1-\alpha)\mathcal{U}(\mu_x, s).
\end{aligned} \tag{23}$$

Let $U(\mu_x, s) = \frac{1}{\pi(s)}\mathcal{U}(\mu_x, s)$. Then the RHS of (23) equals

$$\pi(s)[\alpha U(T_s(\mu_q)) + (1-\alpha)U(\mu_x, s)] + \sum_{s' \neq s} \pi(s')U(T_{s'}(\mu_q)). \tag{24}$$

By property (6), $T(\alpha\mu_q + (1-\alpha)\mu_x) = \alpha T(\mu_q) + (1-\alpha)T(\mu_x)$. Therefore, $T_{s'}(\alpha\mu_q+(1-\alpha)\mu_x) = T(\mu_q)$ for $s' \neq s$, and $T(\alpha\mu_q+(1-\alpha)\mu_x)(s) = \alpha q+(1-\alpha)x$. The latter is a lottery with support in $\Delta(\widehat{C})$, assigning probability $(1-\alpha)$ to the objective unknown consequence x. The result in (8) follows.

Uniqueness of u follows from the uniqueness of U in Lemma 3. By Archemedean continuity, which follows from Axiom 2, for any mixed acts, $\mu \succ \lambda \succ \eta$, there exists $\alpha \in [0,1]$ such that $\lambda \sim \alpha\mu + (1-\alpha)\eta$. This gives uniqueness of $u_x(s, \mu')$ given uniqueness of U.

The uniqueness of π and the fact that $\pi(s) = 0$ if and only if s is null follow from Lemma 3. ■

References

[1] Alon, Shiri (2015): "Worst-case expected utility," *Journal of Mathematical Economics* 60, 43–48.

[2] Anscombe, Francis J. and Robert J. Aumann (1963): "A definition of subjective probability," *Annals of Mathematical Statistics* 34, 199–205.

[3] Blume, Lawrence, David Easley and Joseph Halpern (2021): "Constructive decision theory," *Journal of Economic Theory* 196, 105306.

[4] Board, Oliver and Kim-Sau Chung (2009): "Object-based unawareness: Axioms," Working Paper, University of Minnesota.

[5] Debreu, Gerard (1959): "Topological methods in cardinal utility theory," *Cowles Foundation Discussion Papers* 299.

[6] Dekel, Eddie, Barton L. Lipman and Aldo Rustichini (2001): "Representing preference with a unique subjective state space," *Econometrica* 69, 891–934.

[7] Dominiak, Adam and Gerelt Tserenjigmid (2018): "Belief consistency and invariant risk preferences," *Journal of Mathematical Economics* 79, 157–162.

[8] Dominiak, Adam and Gerelt Tserenjigmid (2022): "Ambiguity under growing awareness," *Journal of Economic Theory* 199, 105256.

[9] Grant, Simon and John Quiggin (2013a): "Inductive reasoning about unawareness," *Economic Theory* 54, 717–755.

[10] Grant, Simon and John Quiggin (2013b): "Bounded awareness, heuristics and the precautionary principle," *Journal of Economic Behavior and Organization* 93, 17–31.

[11] Grant, Simon and John Quiggin (2015): "A preference model for choice subject to surprise," *Theory and Decision* 79, 167–180.

[12] Halpern, Joseph Y. and Leandro C. Rego (2009): "Reasoning about knowledge of unawareness," *Games and Economic Behavior* 67, 503–525.

[13] Halpern, Joseph Y. and Leandro C. Rego (2013): "Reasoning about knowledge of unawareness revisited," *Mathematical Social Sciences* 65, 73–84.

[14] Karni, Edi, Quitze Valenzuela-Stookey, and Marie-Louise Vierø (2021): "Reverse Bayesianism: A generalization," *The B. E. Journal of Theoretical Economics* 21, 557–569.

[15] Karni, Edi and Marie-Louise Vierø (2013): ""Reverse Bayesianism": A choice-based theory of growing awareness," *American Economic Review* 103, 2790–2810.

[16] Karni, Edi and Marie-Louise Vierø (2015): "Probabilistic sophistication and reverse Bayesianism," *Journal of Risk and Uncertainty* 50, 189–208.

[17] Karni, Edi and Marie-Louise Vierø (2017): "Awareness of unawareness: A theory of decision making in the face of ignorance," *Journal of Economic Theory* 168, 301–328.

[18] Kochov, Asen (2018): "A behavioral definition of unforeseen contingencies," *Journal of Economic Theory* 175, 265–290.

[19] Kreps, David M. (1979): "A representation theorem for "preference for flexibility"," *Econometrica* 47, 565–577.

[20] Kreps, David M. (1992): “Static choice in the presence of unforeseen contingencies,” in *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, MIT Press, Cambridge.

[21] Lipman, Barton L. (1999): “Decision theory without logical omniscience: Toward an axiomatic framework for bounded rationality,” *Review of Economic Studies* 66, 339–361.

[22] Ozbay, Erkut Y. (2007): “Unawareness and strategic announcements in games with uncertainty,” in Samet, Dov (Ed.), Proceedings of the Theoretical Aspects of Rationality and Knowledge XI. Presses Universitaires de Louvain, 231–238.

[23] Piermont, Evan (2017): “Introspective unawareness and observable choice,” *Games and Economic Behavior* 106, 134–152.

[24] Schipper, Burkhard C. (2015): “Awareness,” van Ditmarsch, Hans V., Halpern, Joseph Y., van der Hoek, Wiebe, Kooi, Barteld P. (Eds.), Handbook of Epistemic Logic. College Publications, London.

[25] Schipper, Burkhard C. (2022): “Predicting the unpredictable under subjective expected utility,” working paper, UC Davis.

[26] Skiadas, Costis (1997a): “Conditioning and aggregation of preferences,” *Econometrica* 65, 347–367.

[27] Skiadas, Costis (1997b): “Subjective probability under additive aggregation of conditional preferences,” *Journal of Economic Theory* 76, 242–271.

[28] Uyanik, Metin and M. Ali Khan (2022): “The continuity postulate in economic theory: A deconstruction and an integration,” *Journal of Mathematical Economics* 101, 102704.

[29] Vierø, Marie-Louise (2021): “An intertemporal model of growing awareness,” *Journal of Economic Theory* 197, 105351.

[30] Walker, Oliver (2014): “Unawareness with “possible” possible worlds,” *Mathematical Social Sciences* 70, 23–33.

[31] Walker, Oliver and Simon Dietz (2011): "A representation result for choice under conscious unawareness," Working Paper, London School of Economics and Political Science.

Economics Working Papers

2021-05 Giovanni Pellegrino, Efrem Castelnuovo and Giovanni Caggiano: Uncertainty and Monetary Policy during the Great Recession

2021-06 Anna Folke Larsen and Marianne Simonsen: Social emotional learning in the classroom: One-year follow-up results from PERSPEKT 2.0

2021-07 Johannes Schünemann, Holger Strulik and Timo Trimborn: Optimal Demand for Medical and Long-Term Care

2021-08 Tom Lane, Daniele Nosenzo and Silvia Sonderegger: Law and Norms: Empirical Evidence

2021-09 Ina C. Jäkel: Export Credit Guarantees: Direct Effects on the Treated and Spillovers to their Suppliers

2021-10 Martin Paldam: Measuring Democracy - Eight indices: Polity, Freedom House and V-Dem

2021-11 Nicola Maaser and Thomas Stratmann: Costly Voting in Weighted Committees: The case of moral costs

2021-12 Martin M. Andreasen, Giovanni Caggiano, Efrem Castelnuovo and Giovanni Pellegrino: Why Does Risk Matter More in Recessions than in Expansions?

2021-13 Søren Albeck Nielsen: How to Cope with Dyslexia: The Effects of Special Education on Academic Performance, Personality Traits, and Well-being

2022-01 Michael Koch and Ilya Manuylov: Measuring the Technological Bias of Robot Adoption and its Implications for the Aggregate Labor Share

2022-02 Olga Balakina, Claes Bäckman, Andreas Hackethal, Tobin Hanspal and Dominique M. Lammer: Good Peers, Good Apples? Peer Effects in Portfolio Quality

2022-03 Kazuhiko Sumiya and Jesper Bagger: Income Taxes, Gross Hourly Wages, and the Anatomy of Behavioral Responses: Evidence from a Danish Tax Reform

2022-04 Hans Schytte Sigaard: Labor Supply Responsiveness to Tax Reforms

2022-05 Edi Karni and Marie-Louise Vierø: Comparative Incompleteness: Measurement, Behavioral Manifestations and Elicitation

2022-06 Marie-Louise Vierø: Lost in objective translation: Awareness of unawareness when unknowns are not simply unknowns