INSEAD The Business School for the World®

Working Paper

2024/61/DSC
(Revised version of 2024/01/DSC)

Verifiable Uncertainty

Jingyuan Li
Lingnan University, jingyuanli@ln.edu.hk

Ilia Tsetlin
INSEAD, ilia.tsetlin@insead.edu

Fan Wang
ESSEC, Singapore. b00786349@essec.edu

Real-world uncertainties can be difficult to classify into the neat boxes of risk and ambiguity. We suggest verifiability as a practical test for risk, and verifiable uncertainty — often in the form of mechanically implemented equally likely prizes — as a real-world representation of risk. We propose a set of rules by which an agent contrasts a general (unverifiable) event against the verifiable *m* out of *n* chances of winning a prize. Collectively the rules advocate a decision procedure in the spirit of the smooth model of ambiguity, where the agent evaluates verifiably equally likely prizes — an object we call classical lottery — by average (verifiable) utility and, for general prospect, combines various scenarios of how unverifiable uncertainty might unfold. The agent can hold distinct attitudes toward verifiable uncertainty and unverifiable uncertainty. In particular, combining scenarios in a conservative manner reflects aversion to unverifiable uncertainty. We illustrate verifiability as a useful concept for both normative and descriptive decision making.

Key words: Decision Making; Risk; Ambiguity; Uncertainty

Verifiable Uncertainty

Jingyuan Li[*], Ilia Tsetlin[†], Fan Wang[‡]

Abstract

Real-world uncertainties can be difficult to classify into the neat boxes of risk and ambiguity. We suggest verifiability as a practical test for risk, and verifiable uncertainty — often in the form of mechanically implemented equally likely prizes — as a real-world representation of risk. We propose a set of rules by which an agent contrasts a general (unverifiable) event against the verifiable m out of n chances of winning a prize. Collectively the rules advocate a decision procedure in the spirit of the smooth model of ambiguity, where the agent evaluates verifiably equally likely prizes — an object we call classical lottery — by average (verifiable) utility and, for general prospect, combines various scenarios of how unverifiable uncertainty might unfold. The agent can hold distinct attitudes toward verifiable uncertainty and unverifiable uncertainty. In particular, combining scenarios in a conservative manner reflects aversion to unverifiable uncertainty. We illustrate verifiability as a useful concept for both normative and descriptive decision making.

1 Introduction

From vaccine effectiveness to venture capital funding and to nuclear threat, situations abound where precise probabilities are missing. A conscientious agent might run trials, analyze relevant data, draw analogy to historical cases, consult with experts, and the list of options goes on. While taking into account the degrees of evidence that often vary greatly, the agent nonetheless might come to terms with plain estimates like "is the success rate 1 out of 5 or 1 out of 20". Indeed, there seems to be a natural tendency to comprehend uncertainty as frequencies in an equiprobable space. This was evident even for early conceptions of probability theory, as they regularly build on the idea of *classical lottery,* which in its basic form is a finite collection of equally likely prizes.[1] Moreover, in the real world, the popularity of classical-lottery-type devices like Powerball and the game of roulette attests to its effectiveness in communicating randomness to the public.

[*]Lingnan University, Hong Kong. jingyuanli@ln.edu.hk
[†]INSEAD, Singapore. ilia.tsetlin@insead.edu
[‡]ESSEC, Singapore. b00786349@essec.edu

[1]To quote Laplace (1902), "The theory of chance consists in reducing all the events of the same kind to a certain number of cases equally possible."

Classical lotteries implemented by simple mechanics have a special appeal: the kind of uncertainty therein is *verifiable*. People can observe the lottery taking place and the uncertainty getting resolved, with the neat choreography of mechanical instruments to convince them there is nothing to hide. In addition, any stakeholder can potentially check and experiment with the lottery, all the relevant information is laid bare, no subjective inputs are required, there is no room for correlation with any other sources of uncertainty, and, even if the stakeholder ends up with only one prize, she can be assured that all other prizes are actually there. In this sense, verifiable uncertainty implies no uncertainty about uncertainty, and the equiprobability of classical lottery is derived from sufficient information, the opposite to the principle of insufficient reason. In fact, to Keynes and Knight, both forerunners of the field, choosing between classical lotteries is basic and akin to decision under *certainty*.[2]

Verifiability clarifies the meaning of *risk* in practice: uncertainty is known when it is verifiable. Indeed, the classical dichotomy between risk (i.e., objectively known uncertainty, usually described by a known/fixed probability distribution) and *ambiguity* (i.e., unknown probability distribution) still lacks an actionable test on real uncertainties, and here we believe the verifiability criterion comes particularly useful. For example, a projected distribution of global temperature rise is verifiable when seen as a computer simulation. But, unless the model coded in the computer program is taken as literally true, the projection is still unverifiable when seen as a prediction, because only one future is realized and, moreover, how much the future would resemble the past is very much a subjective matter. The same point can be made for actuarial and financial predictions, despite them often being backed by big data sets.

In this paper, we operationalize verifiable uncertainty by substituting it for the risk part of Anscombe and Aumann's (1963) framework, leaving anything unverifiable to ambiguity. Specifically, verifiable uncertainty is modeled via classical lottery, defined as any finite collection of prizes with the understanding that they are equally likely. Classical lottery shall be evaluated by average verifiable utility, and for any prospect that involves unverifiable uncertainty — modeled via states of the world — the agent finds a classical lottery she feels indifferent to and evaluates the prospect the same as that classical lottery. After a brief normative justification of the average verifiable utility model, we focus on how the agent *should* perform the task of finding the indifferent classical lottery. In the typical case of comprehending uncertainty as having m out of n chances to win a prize, for example, a prospect with more good scenarios should be associated with a higher ratio m/n, conveying the idea of monotonicity. Collectively the desiderata we put forward advocate a decision process in the spirit of the smooth model of ambiguity (Klibanoff, Marinacci, and Mukerji 2005, and also Nau 2006), where the agent combines various scenarios of how unverifiable uncertainty may be reduced to verifiable uncertainty. An agent who does this

[2]To quote Keynes (1937), "The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn." A similar view was also expressed in Knight (1921).

in a "conservative" manner exhibits aversion to unverifiable uncertainty.

The agent can have different attitudes towards verifiable and unverifiable uncertainty, and they are separately captured by two utility functions in our model. It is common in the literature to justify such differential treatment based on whether the probabilities are known. Verifiability, by insisting on a practical definition of what is known, makes the contrast of the two types of uncertainty even more stark. In fact, people tend to understand them quite differently (for a review on communicating uncertainty, see Spiegelhalter, Pearson, and Short 2011). And, normatively, verifiable uncertainty is transparent and fosters trust, whereas unverifiable uncertainty is opaque and invites doubt. For example, if a doctor tells a patient that his surgery has a success rate 1/6, that 1/6 might feel quite differently than betting on the outcome of a die toss. The first 1/6 is unverifiable because it depends on how similar the agent is to those in the medical records, and, moreover, the surgery is to be performed on that particular patient only once. Similarly, a vaccine advertised to be 80% chances of being effective for each patient in an i.i.d. manner is different from a vaccine 80% chances effective for all (and 20% chances ineffective for all). The first 80% can be verified in a large population and the advertiser can potentially be held accountable. But policy makers can never be convinced of the second 80%, as the vaccine is to be developed once only.[3] In reality, the differential treatment is ubiquitous, and its profound impact reaches even the vitality of democracy, where the transparency and quality of randomization is essential to political trust and participation (Procaccia 2022, Flanigan et al. 2021, Letsou, Naeh, and Segal 2022).

For a concrete illustration, our model treats the vaccine that is 80% effective in an i.i.d. manner as a (verifiable) classical lottery, where there are 80 boxes out of 100 that contain the effective outcome and the rest 20, the ineffective outcome. The total number of boxes can change as long as there are 80% the good-outcome boxes. This vaccine will be evaluated by average (verifiable) utility, which in this case can be normalized to the proportion of good outcomes, 0.8. In contrast, the vaccine with 80% chances effective for all people is modeled as an act with two states of the world: the "good world" where the vaccine is effective for all (i.e. a classical lottery with all boxes containing the effective outcome) and the "bad world" where the vaccine is ineffective (i.e. all boxes contain the ineffective outcome). The two scenarios are combined through a function $\psi(\cdot)$ to deliver evaluation $\psi^{-1}(0.8\psi(1) + 0.2\psi(0))$. The shape of $\psi(\cdot)$ reflects the agent's attitude towards unverifiable uncertainty — linearity, for example, indicates exchangeability between the two types of uncertainty. Applications as such are further discussed in Section 2.6, where we also contrast the veil of ignorance against other policy uncertainties from the verifiability angle (Rawls 2001, Harsanyi 1955).

[3] We do not make a distinction between ex-ante verifiability, as in examining the mechanics of a roulette wheel or the algorithm that simulates randomness, and ex-post verifiability, as in checking the effectiveness of a vaccine in the population. Running experiments on a randomization device may be for the purpose of familiarizing with the mechanics (i.e., ex-ante verifiability) or collect a sample to approximate the underlying distribution (i.e., ex-post verifiability).

We view the paper's contribution as two-folded. On the conceptual side, we propose verifiable uncertainty as a real-world representation of risk, and we suggest verifiability as a test to clarify which uncertainties are known. Many uncertainties are referred to as risk by professionals, yet according to verifiability, a significant portion of them are not truly known, and a more humble perspective might be called for. Moreover, in light of the critiques pioneered by Allais (1953) and Kahneman and Tversky (1979), restricting risk to verifiable uncertainty may facilitates rational behavior. For example, verifiable small probabilities presented on a large roulette or in a box containing many numbered balls are barely visible and are therefore more likely to be ignored, nudging people away from "anomalies" like the possibility effect. Generally, experimental evidences seem to suggest that people are more comfortable with risk presented as frequencies, and in that context choices tend to be more consistent with expected utility (see, for example, Gigerenzer 1991; 1996, Lipkus and Hollands 1999, and Spiegelhalter et al. 2011).

On the decision-theoretic side, we operationalize verifiable uncertainty by replacing the risk part of the Anscombe-Aumann framework with classical lotteries, and in doing so we give the smooth model of ambiguity a new foundation. Our result on classical lottery mirrors vNM expected utility, and the general model has similar underpinnings as that of classical additive representation. Apart from the lack of a continuum of probability numbers, technical challenges arise because we mostly focus on two-prize prospects, which allow the ideas to be describable in simple terms like "m out of n chances of winning a prize". Our solution should work more broadly. Once classical lotteries are evaluated by average utility, our setting can be seen as a dense version of the Anscombe-Aumann framework, where real-valued probabilities are replaced by rational-valued probabilities. It therefore has the structural richness to host many other models, all with risk replaced by verifiable uncertainty. Detailed discussion and comparison with existing results can be found in Section 2.4.

Section 2 presents an axiomatic treatment of the model, the behavioral content of aversion to unverifiable uncertainty, and an application to social policy. Section 3 concludes the paper. Technical details and proofs are mostly deferred to the Appendix (Section 4).

2 Theory

2.1 Setting

We reformulate the two-stage version of the Anscombe-Aumann framework to take into account verifiable uncertainty. Let X denote a general set of prizes, whose elements can be univariate or multivariate, and quantitative or qualitative. A classical lottery is an I-tuple $\mathbf{x} = (x_1, x_2, ..., x_I)$, with $I \geq 1$. The set of all classical lotteries, $\cup_{I=1}^{\infty} X^I$, is denoted by $\mathcal{X}$. Classical lotteries represent verifiable uncertainty, by which we mean the kind of uncertainty often implemented by simple mechanics, like coin, die, roulette wheel, and computer simulation, for which one can be certain about the mechanism that generates randomness. We do not give verifiability a definition, but we intend its natural meaning in everyday language. In this sense we do not view verifiable uncertainty as always physical; probabilistic statements like “the vaccine is 80% effective in an i.i.d. manner” can be seen to present verifiable uncertainty as well because they can be checked in a population. All the prizes in a classical lottery are deemed equally likely, an intuition we shall carry throughout the paper. The equiprobable assumption may appear limiting, but — because prizes can be repeated in a classical lottery — any rational-valued simple probability distribution will be able to find representative(s) in $\mathcal{X}$.[4]

Uncertainty can also go beyond the transparency of classical lottery and we follow the tradition of Savage (1954) to include *states of the world* into analysis. In doing so we specifically have in mind applications where the agent explicitly considers a set of models (or scenarios) of how uncertainty might unfold: the agent is not sure which model is true, but the randomness within a model is verifiable. A state of the world effectively stands for such a model, and the remaining uncertainty conditioned on a state can be captured by classical lottery. Formally, set $\mathcal{S}$ consists of finitely many states of the world and an *act* is any function that maps states $\mathcal{S}$ to classical lotteries $\mathcal{X}$. For every act f, only one state $s \in \mathcal{S}$ and, correspondingly, only one classical lottery $f(s)$ shall be realized. Hence, even if a probability distribution on $\mathcal{S}$ is given, there is no way to check it. In fact, the agent cannot even verify whether the classical lottery $f(s')$ could have been delivered as promised had state $s' \in \mathcal{S}$ materialized instead. In this sense, acts presents unverifiable uncertainty on top of verifiable uncertainty.

For a concrete example, consider a group of scientists who disagree on models of disease transmission. A state here stands for a fully specified model, conditional on which randomness is mechanically simulated by computer. Correspondingly, a disease containment policy can be thought of as an act, which associates every model with a simulated verifiable distribution of health-related outcomes over the population. This type of model uncertainty has been discussed by Marinacci (2015). Our setting also comes natural when the two types of uncertainty are well defined and independent, as in Anscombe and Aumann's original conception of horse-roulette lottery: each state is a scenario where some horse wins, and the uncertainty in the roulette is mechanical and verifiable. A further discussion about the applicability of our result is deferred to Section 2.6.

It is not an ad hoc choice that unverifiable uncertainty precedes verifiable uncertainty in the modeling of acts. The reversed setting, where a prize in a classical lottery may itself contain unverifiable uncertainty, would render classical lotteries “opaque” and not

[4]We focus on classical lotteries for their natural simplicity, though they may not cover all verifiable uncertainties (we thank a reviewer for making this point). However, even in cases like the probability of a random dot falling inside a circle as opposed to the square exactly containing it, people tend to use discretization — in our language this is a form of classical lottery — as the benchmark of understanding. In general, a probability distribution often admits approximation by empirical distributions, which are verifiable, finite, and rational-valued (e.g., the Glivenko-Cantelli Theorem).

thoroughly inspectable. That is, classical lottery would lose verifiability, which makes the setting conceptually inconsistent. For this reason, verifiability lends to a natural discipline on the otherwise flexible framework, and it foresees the separation of uncertainty in our model.[5]

An agent needs to form a *preference* over the set $\mathcal{F} = \mathcal{X}^{\mathcal{S}}$ of acts, which is modeled as a binary relation $\succcurlyeq$ on $\mathcal{F}$, with its asymmetric part $\succ$ standing for strict preference and the symmetric part $\sim$ standing for indifference. Following a literature convention, a constant act that maps all states to a single classical lottery is identified as that classical lottery. The preference therefore needs to rank non-constant acts against classical lotteries, highlighting the contrast between the two types of uncertainties. Our goal is to introduce a model of arriving at such preferences through a series of normative principles that we deem reasonable. To start with, we require this preference to be a weak order. Recall that a binary relation $\succcurlyeq$ on $\mathcal{F}$ is *complete* if for any two acts $f, g \in \mathcal{F}$, either $f \succcurlyeq g$ or $g \succcurlyeq f$, and it is *transitive* if for any three acts $f, g, h \in \mathcal{F}$, $f \succcurlyeq g$ and $g \succcurlyeq h$ imply $f \succcurlyeq h$.

Axiom 1 (Weak Order): The preference $\succcurlyeq$ is complete and transitive.

From the normative standpoint, completeness can be desirable because a call for decision may arise between any two acts, and transitivity is a standard tenet of rationality.

2.2 Evaluating Classical Lottery

We shall propose two principles for evaluating classical lotteries, all based on the idea that the prizes are equally likely. The following notations are adopted to streamline the exposition. For any $\mathbf{x} \in \mathcal{X}$, $|\mathbf{x}|$ denotes the length of $\mathbf{x}$. For any two classical lotteries $\mathbf{x} = (x_1, x_2, ..., x_k)$ and $\mathbf{y} = (y_1, y_2, ..., y_j)$, the concatenation of $\mathbf{x}$ and $\mathbf{y}$, denoted by $\mathbf{x} \vee \mathbf{y}$, is the $(k + j)$-tuple $(x_1, x_2, ..., x_k, y_1, y_2, ..., y_j)$. For any $\mathbf{x} \in \mathcal{X}$, let $\mathbf{x}^n = \mathbf{x}^{n-1} \vee \mathbf{x}$, $n \geq 2$. For any classical lottery $\mathbf{x} \in \mathcal{X}$ and any set $E \subset \{1, 2, ..., |\mathbf{x}|\}$, $\mathbf{x}_E$ denotes the classical lottery $(x_i)_{i \in E}$ where i is ordered from the smallest to the largest.

To motivate the two principles, let us first consider what it means for the prizes to be equally likely. To begin with, it implies that any permutation on a classical lottery should be immaterial. Formally, $\mathbf{x} \vee \mathbf{y} \sim \mathbf{y} \vee \mathbf{x}$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$. Moreover, it also demands that any replication of a classical lottery should not matter, that is, $\mathbf{x}^n \sim \mathbf{x}$ for all $\mathbf{x} \in \mathcal{X}$. The idea is that, as the prizes are equally likely, $\mathbf{x}^n$ can be seen as n copies of the same classical lottery $\mathbf{x}$, and therefore it is $\mathbf{x}$ that ultimately matters. These two properties have indeed

[5] For a recent example on the framework's flexibility, see He (2021), where "difficult uncertainties" follow "simple uncertainties". As to multi-stage lotteries, which we do not touch on in this paper, a setting conceptually consistent to our motivation would be for the lottery to resolve unverifiable uncertainty up to some stage, after which there can only be verifiable uncertainty. For example, the Galton box — a device used to demonstrate the central limit theorem — can be seen as a multi-stage classical lottery, where each stage has two outcomes that are verifiably equally likely. Finally, see Li, Rohde, and Wakker (2023) for a discussion of (partial) separability.

been adopted in the literature (see, for example, Kothiyal, Spinu, and Wakker 2014), and they are implicitly invoked in formulating our principles albeit not directly assumed.

The first principle enables the agent to simplify any comparison between classical lotteries by cancelling out their common parts. For example, if a prize appears 1 out of 2 times in classical lottery $\mathbf{x}$ and 2 out of 4 times in classical lottery $\mathbf{y}$, then in both cases this prize obtains with the same likelihood of 50% and hence it should not matter for the relative ranking of $\mathbf{x}$ and $\mathbf{y}$. In general, if a prize obtains with the same frequency in two classical lotteries, then it should not matter for their relative ranking. This principle can be formally stated as follows.

Axiom 2 (Cancellation): For all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, $E \subset \{1, 2, ..., |\mathbf{x}|\}$, and $F \subset \{1, 2, ..., |\mathbf{y}|\}$, if $\frac{|E|}{|\mathbf{x}|} = \frac{|F|}{|\mathbf{y}|}$ and $\mathbf{x}(i) = \mathbf{y}(j)$ for all $i \in E$ and $j \in F$, then $\mathbf{x} \succcurlyeq \mathbf{y}$ if and only if $\mathbf{x}_{E^c} \succcurlyeq \mathbf{y}_{F^c}$.

The second principle requires the relative ranking of two classical lotteries not to reverse when one of them gets perturbed just a little. Specifically, for any classical lottery $\mathbf{x}$ and $\mathbf{z}$, the concatenation $\mathbf{x} \vee \mathbf{z}$ might present a major modification on $\mathbf{x}$, especially when the length of $\mathbf{z}$ is large relative to that of $\mathbf{x}$. However, since prizes are equally likely, $\mathbf{x}^n$ can be understood as a bigger version of $\mathbf{x}$, and hence $\mathbf{x}^n \vee \mathbf{z}$, a smaller modification on $\mathbf{x}$. The principle can therefore be stated as follows.

Axiom 3 (Archimedeanity): For all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathcal{X}$ with $\mathbf{x} \succ \mathbf{y}$, there exists an $N \in \mathbb{N}$ such that for all $n \geq N$, $(\mathbf{x}^n \vee \mathbf{z}) \succ \mathbf{y}$ and $\mathbf{x} \succ (\mathbf{y}^n \vee \mathbf{z})$.

These two principles can be seen as the classical-lottery version of the independence and Archimedeanity axioms of the expected utility theory of von Neumann and Morgenstern (1944, hereafter vNM).[6] And, as the following result shows, they are necessary and sufficient for a weak-order preference to be representable by average (verifiable) utility, which corresponds to expected utility for classical lottery.

Proposition 1. *A weak-order preference* $\succcurlyeq$ *on* $\mathcal{X}$ *satisfies Cancellation and Archimedeanity if and only if there is a utility function* $u : X \to \mathbb{R}$ *such that for all* $\mathbf{x}, \mathbf{y} \in \mathcal{X}$*,*

$$\mathbf{x} \succcurlyeq \mathbf{y} \Leftrightarrow \frac{1}{|\mathbf{x}|} \sum_{i=1}^{|\mathbf{x}|} u(x_i) \geq \frac{1}{|\mathbf{y}|} \sum_{j=1}^{|\mathbf{y}|} u(y_j). \tag{2.1}$$

Moreover, the utility function $u(\cdot)$ *is unique up to affine transformation.*

The proof hinges on the same vNM strategy of simplifying a classical lottery gradually into a two-prize classical lottery, where the evaluation can be shown to be decided by the

[6] The vNM Archimedeanity requires any strict preference to be robust to *some* small perturbation. Consecutive applications of the condition imply robustness up to infinitely many small perturbations. However, this is still technically weaker than our Archimedeanity, which calls for *all* the small perturbations associated with large enough n. The conveyed normative idea remains the same — small perturbations should barely matter.

frequency of the better prize. In comparison with existing results on average utility (for example, Kothiyal, Spinu, and Wakker 2014), our axioms and techniques display a close proximity with vNM's axiomatization. Proposition 1 can be seen as a direct counterpart of vNM expected utility in a simpler domain. Moreover, classical lottery is well suited to von Neumann and Morgenstern's frequency interpretation of probability.

For ease of exposition, we shall, for any utility function u over prizes, adopt the corresponding capital-letter function $U : \mathbf{x} \mapsto \frac{1}{|\mathbf{x}|}\sum_{i=1}^{|\mathbf{x}|} u(x_i)$ to evaluate classical lotteries. Function U will be referred to as *average verifiable utility*.

2.3 Classical Lottery to Gauge Unverifiable Uncertainty

Acts generally involve unverifiable uncertainty, which cannot be inspected or objectively quantified. To evaluate acts, we instead start with the simple principle that the preference over classical lotteries should be carried over to any state.

> **Axiom 4 (Monotonicity)**: For any two acts $f, g \in \mathcal{F}$, if $f(s) \succcurlyeq g(s)$ for all $s \in \mathcal{S}$, then $f \succcurlyeq g$.

We now propose a procedure of evaluating two-prize acts. The simplicity of these acts makes it easier to communicate normative principles, and they are also rich enough to approximate general acts in terms of desirability. To streamline the notation, for any two prizes $x \succ y$, let $\mathcal{X}_{x,y}$ denote the set of classical lotteries that involve only the two prizes x and y, and let $\mathcal{F}_{x,y}$ denote the set of acts that maps states $\mathcal{S}$ to $\mathcal{X}_{x,y}$. For any classical lottery $\mathbf{x} \in \mathcal{X}_{x,y}$, let $r_x(\mathbf{x})$ denote the (relative) frequency of prize x in the lottery, namely, $r_x(\mathbf{x}) = |\{i : \mathbf{x}(i) = x\}|/|\mathbf{x}|$. Under Axioms 1–4, any two-prize act $f = (f(s))_{s\in\mathcal{S}} \in \mathcal{F}_{x,y}$ can be seen as a profile $(r_x(f(s)))_{s\in\mathcal{S}}$ of state-contingent frequencies of the better prize x. The crucial step of the procedure is to aggregate this profile into a single frequency, represented by a *matching classical lottery* $\mathbf{x}_f \in \mathcal{X}_{x,y}$ that the agent feels indifferent to. The act can then be evaluated by the average utility of $\mathbf{x}_f$. This procedure can be summarized as follows,

$$f \in \mathcal{F}_{x,y} \;\sim\; \mathbf{x}_f \in \mathcal{X}_{x,y} \;\mapsto\; U(\mathbf{x}_f). \tag{2.2}$$

Since the matching classical lottery $\mathbf{x}_f$ is generally not a sure prize, this procedure is different from the conventional one in the literature where the agent tries to find the act's certainty equivalent (or willingness to pay). In a nutshell, the procedure reduces the unverifiable uncertainty in f to the verifiable uncertainty in $\mathbf{x}_f$. In terms of real behavior, this can entail considering and combining various scenarios of how unverifiable uncertainty might unfold. A similar procedure was used by Wang (2019) to motivate a definition of comparative ambiguity attitudes.

Setting aside the normative content, there are two technical issues surrounding procedure (2.2). First, as frequency only takes value in rational numbers, some act f may not have a matching classical lottery $\mathbf{x}_f$. We can circumvent this issue by considering the set of classical lotteries that are less preferred than f, $\{\mathbf{x} \in \mathcal{X}_{x,y} : f \succcurlyeq \mathbf{x}\}$, or its opposite, $\{\mathbf{x} \in \mathcal{X}_{x,y} : \mathbf{x} \succcurlyeq f\}$. For simplicity we characterize these sets by their mathematical limits, which give us a notion of *matching frequency*.

Definition 2. Under Axioms 1–4, for any two prizes $x \succ y$ and any act $f \in \mathcal{F}_{x,y}$, the matching frequency of f is

$$m_{x,y}(f) = sup\{r_x(\mathbf{x}) : f \succcurlyeq \mathbf{x}, \mathbf{x} \in \mathcal{X}_{x,y}\}.$$

Since $m_{x,y}(f)$ may not be a rational number, it does not necessarily have a representing classical lottery $\mathbf{x}_f \in \mathcal{X}_{x,y}$. Instead, it refers to the set $\{\mathbf{x} \in \mathcal{X}_{x,y} : f \succcurlyeq \mathbf{x}\}$ of classical lotteries in $\mathcal{X}_{x,y}$ that bounds the desirability of act $f \in \mathcal{F}_{x,y}$. Correspondingly, the behavioral meaning of $m_{x,y}(f) \geq m_{x,y}(g)$ is that, although the agent may have difficulty comparing acts $f, g \in \mathcal{F}_{x,y}$ directly, she is able to reach the following conclusion: for any $\mathbf{x} \in \mathcal{X}_{x,y}$, $f \succcurlyeq \mathbf{x}$ whenever $g \succcurlyeq \mathbf{x}$. On a side note, matching frequency can be equivalently defined as $inf\{r_x(\mathbf{x}) : \mathbf{x} \succcurlyeq f, \mathbf{x} \in \mathcal{X}_{x,y}\}$.

The second technical issue concerns whether classical lotteries are rich enough to effectively gauge general uncertainty. We therefore impose the following axiom that the preference should allow classical lotteries to "tell acts apart".

> **Axiom 5 (Denseness)**: For any two acts $f \succ g$, there exists a classical lottery $\mathbf{x}$ such that $f \succ \mathbf{x} \succ g$.

As the following result shows, under Denseness, matching frequency fully captures the preference relation on binary acts. In the Appendix we discuss an example where Axioms 1–4 are satisfied but Denseness is not.

Lemma 3. *Under Axioms 1–5, for any two prizes $x \succ y$ and any two acts $f, g \in \mathcal{F}_{x,y}$,*

$$f \succcurlyeq g \Longleftrightarrow m_{x,y}(f) \geq m_{x,y}(g).$$

Matching frequency $m_{x,y}(\cdot)$ on the domain $\mathcal{F}_{x,y}$ of binary acts admits a natural replica on the domain $\mathbb{Q}_{[0,1]}^{\mathcal{S}}$ of state-contingent frequencies. Here $\mathbb{Q}_{[0,1]}$ stands for the set of rational numbers in $[0, 1]$, and a vector $\mathbf{q} = (q_s)_{s \in \mathcal{S}} \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$ stands for the state-contingent frequency q_s of prize x (and frequency $1 - q_s$ of prize y) for all $s \in \mathcal{S}$. Indeed, any $\mathbf{q} \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$ inherits a matching frequency — denoted by $m_{x,y}(\mathbf{q})$ — from any of its *representing binary act*, namely, $m_{x,y}(\mathbf{q}) \equiv m_{x,y}(f)$ for any $f \in \mathcal{F}_{x,y}$ with $\mathbf{q}(s) = r_x(f(s))$ for all $s \in \mathcal{S}$. Under Axioms 1–4, any two such representing binary acts should be indifferent.

We shall propose two desiderata to discipline the procedure (2.2), or, equivalently, the calculation of $m_{x,y}(\mathbf{q})$. Together they pin down model (2.4), where state-contingent frequency $\mathbf{q}$ is aggregated into matching frequency $m_{x,y}(\mathbf{q})$ in the style of subjective expected utility, and the utility itself reflects attitudes toward unverifiable uncertainty. The first desideratum requires $m_{x,y}(\mathbf{q})$ to be continuous on $\mathbb{Q}_{[0,1]}^{\mathcal{S}}$.

Axiom 6 (Continuity): For any two prizes $x \succ y$ and any convergent vectors $\{\mathbf{q}_n\}_{n=1}^{\infty} \subset \mathbb{Q}_{[0,1]}^{\mathcal{S}}$, $\{m_{x,y}(\mathbf{q}_n)\}_{n=1}^{\infty}$ should also be convergent.[7]

The axiom can be alternatively stated in terms of preference: for any $a > b$ in $\mathbb{Q}_{[0,1]}$, with acts $f, g, h_n \in \mathcal{F}_{x,y}$ representing $\mathbf{a}, \mathbf{b}$, and $\mathbf{q}_n$, respectively, $n \geq 1$, it cannot be that $h_n \not\succcurlyeq f$ infinitely often and also $g \not\succcurlyeq h_n$ infinitely often. Here and after, the bold version of a number stands for the constant vector of that number.

The second desideratum conveys the key behavioral content of our model. It essentially requires frequencies on different states to be taken into account separately but in the same scale. To explain how $m_{x,y}(\cdot)$ contains an evaluation of frequencies, we adopt the notation that, for any state $s \in \mathcal{S}$, number $a \in \mathbb{Q}_{[0,1]}$, and vector $\mathbf{q} \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$, $a_s\mathbf{q}$ stands for the vector that assumes value a on state s and value $\mathbf{q}(s')$ for all state $s' \neq s$. For any vectors $\mathbf{p}, \mathbf{q} \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$ and $b, b+\epsilon \in \mathbb{Q}_{[0,1]}$, $m_{x,y}(b_s\mathbf{q}) \geq m_{x,y}((b+\epsilon)_s\mathbf{p})$ indicates that the frequency incremental from b to $b+\epsilon$ in state s is not enough to offset the disadvantage of $\mathbf{p}$ relative to $\mathbf{q}$ outside s. If we similarly observe $m_{x,y}((a+\delta)_s\mathbf{p}) \geq m_{x,y}(a_s\mathbf{q})$, which shows that the same disadvantage is overcome, then a δ incremental on a seems more significant than an ϵ incremental on b. For this line of reasoning to be valid, the state s should matter (i.e., not a zero-probability state), and moreover, the conclusion should not be contradictory were the benchmark — the $\mathbf{p}$ and $\mathbf{q}$ difference outside s — to change. The desideratum is a formal statement of these conditions. Notice that, for any prize pair $x \succ y$, a state $s \in \mathcal{S}$ is *essential* if $m_{x,y}(1_s\mathbf{q}) > m_{x,y}(0_s\mathbf{q})$ for some vector $\mathbf{q} \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$.

Axiom 7 (Consistent Reduction): For any prize pair $x \succ y$, essential states s and s', frequencies $a, a+\delta, b, b+\epsilon \in \mathbb{Q}_{[0,1]}$, and vectors $\mathbf{p}, \mathbf{q}, \mathbf{p}', \mathbf{q}' \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$, if we have $m_{x,y}(b_s\mathbf{q}) \geq m_{x,y}((b+\epsilon)_s\mathbf{p})$ and $m_{x,y}((a+\delta)_s\mathbf{p}) \geq m_{x,y}(a_s\mathbf{q})$, then $m_{x,y}(a_{s'}\mathbf{q}') \geq m_{x,y}((a+\delta)_{s'}\mathbf{p}')$ implies $m_{x,y}(b_{s'}\mathbf{q}') \geq m_{x,y}((b+\epsilon)_{s'}\mathbf{p}')$.

This axiom carries the classical idea of consistent tradeoffs (Wakker 1989), though there are two differences here. First, our axiom applies to each prize pair $x \succ y$, and importantly, it does allow the tradeoffs to be different across prize pairs. Second, the domain we have for each prize pair is $\mathbb{Q}_{[0,1]}^{\mathcal{S}}$, while the classical idea is formulated on connected topology. It is tempting to use the cardinal value of matching frequency to measure the desirability of frequencies. For example, $m_{x,y}((a+\epsilon)_s\mathbf{p}) - m_{x,y}(a_s\mathbf{p})$ being larger than $m_{x,y}((b+\epsilon)_s\mathbf{p}) - m_{x,y}(b_s\mathbf{p})$ also seems to indicate that a δ incremental on a is more significant than an ϵ incremental on b. But this logic ignores the tradeoffs across states, which is basically what attitude toward unverifiable uncertainty is about. Indeed, if a on state s is small in comparison with frequencies outside s, a cautious agent might worry about s happening and therefore value especially an incremental on a. That is, desirability measured as such

[7] We use the standard notion of convergence according to Euclidean distance in $\mathbb{R}^n$. Axiom 6 is stronger than the usual continuity condition (i.e., for any $\mathbf{q} \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$ and $\{\mathbf{q}_n\}_{n=1}^{\infty} \subset \mathbb{Q}_{[0,1]}^{\mathcal{S}}$, $m_{x,y}(\mathbf{q}_n)$ converges to $m_{x,y}(\mathbf{q})$ whenever $\mathbf{q}_n$ converges to $\mathbf{q}$), which still allows $m_{x,y}(\mathbf{q}_n)$ to diverge when $\mathbf{q}_n$ converges to a non-rational vector.

is dependent on $\mathbf{p}$. A similar problem appears in the calculation of willingness to pay for deterministic multivariate object (Smith and Dyer, 2021).

So far the axioms have been mostly about evaluating two-prize acts. However, the Monotonicity axiom allows us to use the evaluation of two-prize acts to approximate that of general acts.[8] Specifically, by this axiom, a general act f should be indifferent to a two-prize act g if $f(s) \sim g(s)$ for all s, and we can then use the valuation of g for f as well. Due to the structure of classical lotteries, we may not always be able to find such an indifferent two-prize act, and in these cases we will need to approximate the act f using two-prize acts that are, for example, supported on the worst and the best prizes involved by f.

2.4 The Main Result

We can now state our main result, which establishes a model of decision under general uncertainty. Notice that, for any set E of real numbers, $conv(E)$ denotes the convex hull of E.

Theorem 4. *Preference relation $\succcurlyeq$ satisfies Axioms 1–7 if and only if $\succcurlyeq$ admits the representation that*

$$f \succcurlyeq g \Leftrightarrow \int_{\mathcal{S}} \psi(U(f(s)))dP \geq \int_{\mathcal{S}} \psi(U(g(s)))dP, \tag{2.3}$$

where $U : \mathcal{X} \to \mathbb{R}$ is average verifiable utility, $\psi : conv(U(\mathcal{X})) \to \mathbb{R}$ is a continuous, strictly increasing function, and P is a probability distribution over $\mathcal{S}$. Furthermore, P is unique when there exists two prizes x and y such that $x \succ y$. And, when there are at least two essential states, ψ is unique up to affine transformation given U.

The function U is the average verifiable utility characterized in Proposition 1.[9] The probability distribution P encapsulates the agent's subjective probability assessments over the state space, as in Savage (1954). The function ψ captures the agent's attitudes toward unverifiable uncertainty. In particular, when ψ is affine, the agent linearly aggregates state-contingent utilities according to the weights mandated by P, as if the two types of uncertainty are interchangeable. Non-linear ψ then reflects the agent's distinct attitudes toward unverifiable uncertainty, a topic we cover in Section 2.5. Finally, since set $U(\mathcal{X})$ — the utilities of classical lotteries — can be a non-connected set of real numbers, we instead define ψ on $conv(U(\mathcal{X}))$ for an easy statement of the continuity property.

[8] For a side note, we would like to point out that Axiom 7 actually implies Axiom 4 but only restricted to two-prize acts. That is, for all classical lotteries $\mathbf{x}, \mathbf{y} \in \mathcal{X}_{x,y}$, all act $f \in \mathcal{F}_{x,y}$, and all essential state s, $\mathbf{x} \succ \mathbf{y}$ if and only if $\mathbf{x}_s f \succ \mathbf{y}_s f$. The full Axiom 4 also applies to acts that involve more than two prizes.

[9] Though we use U in the model, any continuous strictly increasing transformation of U (with ψ adjusted correspondingly) is qualified for the role, as it only represents the ranking over classical lotteries. For this reason, we choose not to state a uniqueness property for U in Theorem 4. The function $\psi \circ U$ does carry a cardinal meaning as it weighs state-contingent classical lotteries against each other (Keeney and Raiffa, 1993).

Restricted to two-prize acts, model (2.3) gives a precise expression of the reduction procedure (2.2). Namely, for any prizes $x \succ y$, take any two-prize act $f \in \mathcal{F}_{x,y}$, and its matching frequency $m_{x,y}(f)$ can be written as follows,

$$m_{x,y}(f) = \psi^{-1}\left(\int_{\mathcal{S}} \psi\left(r_x(f(s))\right) dP\right), \tag{2.4}$$

under the normalization $U(x) = 1$ and $U(y) = 0$. Conditional on any state $s \in \mathcal{S}$, verifiable uncertainty as frequency $r_x(f(s))$ is taken into account simply as a cardinal number in this formula. The states $\mathcal{S}$, however, involve unverifiable uncertainty, quantified by subjective probability distribution P, and the model does not treat it interchangeably with frequencies (which would have let to $m_{x,y}(f) = \int_{\mathcal{S}} r_x\left(f(s)\right) dP$). Rather, the frequency profile $(r_x(f(s)))_{s \in \mathcal{S}}$ is aggregated through a function ψ that reflects the attitudes toward unverifiable uncertainty. Concave ψ, for example, implies a discount on (expected) frequency just like risk aversion implies a discount on expected value.

Discussion of the Result

A few comments are provided here to clarify the technical novelty and potential of our result. Foremost, the setting we use closely resembles the Anscombe-Aumann framework, which has been proven a fertile ground for models of decision under uncertainty (e.g., Schmeidler 1989, Gilboa and Schmeidler 1989, and Klibanoff, Marinacci, and Mukerji 2005). However, whereas the standard approach utilizes the powerful tools of functional analysis applied to the linear structure of the framework, we use frequency tradeoffs to deliver model (2.4) for every prize pair $x \succ y$, which is technically closer to classical results on additive representation (Wakker 1989). The challenges are (1) the domains like $\mathcal{F}_{x,y}$ are non-connected and (2) results for different prize pairs need to be brought together to deliver model (2.3). The continuity and the monotonicity axioms are instrumental in resolving these difficulties.

The literature on multi-layered uncertainty and the smooth model of ambiguity in particular is extensive, and our discussion will inevitably be incomplete. The model was proposed by Klibanoff, Marinacci, and Mukerji (2005) in a setting akin to the two-stage Anscombe-Aumann framework, where both risk and ambiguity are exogenously given. They assume as primitives two types of preferences — over first-order and second-order acts, respectively — that satisfy the Savage axioms, and connect them through a monotonicity condition to deliver the model. Seo (2009) gives a more direct axiomatization with a single preference in the larger three-stage Anscombe-Aumann setting, and Denti and Pomatto (2022) clarifies the model's identifiability property in the original two-stage setting. Models in the same spirit also take form in other domains. For example, on the purely objective side are compound lotteries (Kreps and Porteus 1978 and Segal 1987; 1990) and menus of lotteries (Ahn 2008), and on the purely subjective side there are bets on product state space (Nau

2006) and Savage acts (Ergin and Gul 2009). And, He (2021) suggests a "reversed" version of the model where hard uncertainty follows easy uncertainty. In comparison with the literature, the approach we take here is more concrete and application oriented. Indeed, our conception of verifiable uncertainty is driven by how risks are represented in reality, and so is our technical effort to reformulate the smooth model with equiprobable risks. In particular, the procedure (2.2) uses the natural concept of winning chances to make the normative axioms practical, and it imposes no restriction on prizes, which can well be qualitative. Our formulation of states is also practice driven, as they correspond to cases where the decision maker explicitly considers multiple scenarios or models, whereas the literature usually takes states as abstract and often requires infinitely many.

We end the discussion with a note on the potential of our approach. Under Axioms 1–4, classical lotteries can be identified as standard simple lotteries with rational probabilities. In this sense, our setting is a mathematically dense version of the Anscombe-Aumann two-stage setup, and it is therefore rich enough to host many other generalizations. On the classical lottery part, one can for example set out to incorporate the size effect, that is, to treat 1 out of 2 and 50 out of 100 differently. This has been addressed by Flores-Szwagrzak (2024) from an information content perspective, but in cases like income profile the size of a society may be interesting in more direct ways. Another immediate direction is to incorporate psychological tendencies in probability perception in the style of rank-dependent utility (Quiggin 1982, Wang 2022). On the unverifiable uncertainty side, various forms of non-expected utility can be explored. For example, a more radical aversion to unverifiable uncertainty can be captured by an adaption of the max-min model (Gilboa and Schmeidler 1989). Finally, it can be natural and interesting to limit the maximum length of a classical lottery. Normally the model will only be partially identified under a finite preference relation. However, one can hope to achieve unique identification by assuming a bit more. For example, Weber's idea of just-noticeable difference might motivate a preference solvability condition (e.g., given the best and worst prizes $x \succ y$, for all other prize z, $z \sim \mathbf{x}$ for some $\mathbf{x} \in \mathcal{X}_{x,y}$), which effectively restricts the precision of utility function to match the richness of preference. One can also introduce a continuum of prizes and treat the boxes as multiple attributes (Keeney and Raiffa 1993).

2.5 Aversion to Unverifiable Uncertainty

What does it mean, in terms of behavior, for an agent to be averse to unverifiable uncertainty? For any two-prize act $f \in \mathcal{F}_{x,y}$ with $x \succ y$, we may expect this agent to aggregate state-contingent winning probabilities $(r_x(f(s)))_{s\in\mathcal{S}}$ *with a discount.* We formalize this intuition from the hedging angle of Schmeidler (1989) and Gilboa and Schmeidler (1989). To fix the idea, suppose that there are only two states s_1 and s_2. Consider two-prize acts $f, g \in \mathcal{F}_{x,y}$ with $x \succ y$, where f does better in s_1 (i.e., $r_x(f(s_1)) > r_x(f(s_2))$) and g does better in s_2. Aversion to unverifiable uncertainty might mean that the agent worries about s_2

happening when choosing f and, similarly, worries about s_1 when choosing g. Hence, if she finds f and g equally attractive, she would like to use the strength of f to offset the weakness of g, and vice versa. That is, if $h \in \mathcal{F}_{x,y}$ is such that $r_x(h(s_n)) = \frac{1}{2}r_x(f(s_n)) + \frac{1}{2}r_x(g(s_n))$ for $n \in \{1, 2\}$, the agent would (weakly) prefer h to both f and g.

Generally, for any prizes $x \succ y$ and acts $f, g, h \in \mathcal{F}_{x,y}$, act h is said to be a *compromise* between acts f and g if there exists a (rational) number $\alpha \in (0, 1)$ such that $r_x(h(s)) = \alpha r_x(f(s)) + (1 - \alpha)r_x(g(s))$ for all $s \in \mathcal{S}$. When $f \sim g$, the compromise h should appear more attractive to an agent who is averse to unverifiable uncertainty. However, because we do not always have a rich set of indifferent acts in our setting, we require instead that the compromise act h cannot do worse than the lesser act between f and g.

Definition 5. Preference $\succcurlyeq$ exhibits *unverifiable uncertainty aversion* (respectively, *unverifiable uncertainty seeking*) if, for all prizes $x \succ y$ and acts $f, g \in \mathcal{F}_{x,y}$ with $f \succcurlyeq g$, we have $h \succcurlyeq g$ (respectively, $f \succcurlyeq h$) for all $h \in \mathcal{F}_{x,y}$ that is a compromise between f and g.

The next result equates unverifiable uncertainty aversion to the function ψ being concave. The mathematical underpinning is based on Debreu and Koopmans (1982).

Proposition 6. *Suppose that preference $\succcurlyeq$ can be represented by model 2.3 with (P, ψ, U) and that there are at least two essential states. Then, $\succcurlyeq$ exhibits unverifiable uncertainty aversion (respectively, unverifiable uncertainty seeking) if and only if $\psi(\cdot)$ is concave (respectively, convex).*

Since $\psi \circ U$ can be interpreted as a utility function for unverifiable uncertainty, in language of risk aversion, unverifiable uncertainty aversion (i.e., concave ψ) can be understood as the agent being more "risk averse" in the domain of unverifiable uncertainty. This result provides the theoretical foundation for economic analysis from the verifiability angle. For example, consider a principal who commissions an agent to bring in investment opportunities. When a project has verifiable chances of success, the agent can be excused no matter what the outcome is. And for this reason the principal may be happy to pay the agent ex ante (i.e., before the resolution of uncertainty). However, when the chances are unverifiable, the principal — worrying about moral hazard — may exhibit aversion to unverifiable uncertainty and insist on paying ex post. The agent therefore enjoys a sure payoff in the verifiable case while bearing all the uncertainty in the unverifiable case, and this may in turn have implications for project procurement.

2.6 An Illustration

We now connect model (2.3) to the problem of choosing social policy. This is a natural application because, as we argue below, many social policies present both verifiable and unverifiable uncertainty. The vaccine problem we discussed in the introduction is one example. And more generally, we focus on policies that allow for an easy distinction between the outcome of a policy, which describes what each citizen gets as a result of the policy, and the uncertainty of the outcome itself, which comes from the planner's limited knowledge (Harsanyi 1955). In the case of monetary prizes, a policy's outcome is captured by an income profile, modeled as any I-tuple $\mathbf{x} = (x_1, x_2, ..., x_I) \in \mathbb{R}^I$, $I \geq 1$, with the interpretation that person i gets income x_i, $1 \leq i \leq I$. To take into account the uncertainty of outcome, we model social policy as any state-contingent outcome, and in the monetary case this corresponds to a mapping from some state space $\mathcal{S}$ to the set of income profiles $\cup_{I=1}^{\infty}\mathbb{R}^I$. A state $s \in \mathcal{S}$ here should provide enough specification so that one outcome (e.g., one income profile) arises.

The planner faces two types of uncertainty when evaluating a social policy. For a fixed policy outcome, the veil of ignorance approach requires the planner to forget about her identity and imagine that she could become any person of the society (Rawls 2001). In the example of an income profile $(x_1, x_2, ..., x_I)$, the planner is not sure which identity $i \in \{1, 2, ..., I\}$ she will assume and obtain prize x_i. Similarly, in the case of a vaccine that is 80% effective in an i.i.d. manner, the outcome is a classical lottery with 80% citizens getting effective treatment, and without an identity the planner is not sure whether her treatment will be effective or not. Be that as it may, this identity uncertainty is verifiable because, even though she takes only one identity, she can in principle check (from census data or simply by looking around) whether a policy outcome is realized as described. On the other hand, the outcome uncertainty modeled via states of the world is unverifiable, because only one outcome (or one world) realizes and there is no way to check whether the unrealized outcomes are actually plausible, let alone with any specific probabilities. The two types of uncertainty we discuss here echo with the identity and outcome uncertainties of Grant et al. (2010), which in their paper are both treated as lotteries directly. Aside from the verifiability angle, our model's order of uncertainty aggregation — which again is dictated by the verifiability concept — is the opposite to theirs (i.e., they evaluate outcome uncertainty inside identity uncertainty).

According to model (2.3), the planner evaluates a policy outcome by simple utilitarianism.[10] In the vaccine example, this corresponds to the percentage of effective treatment in an outcome (under the normalization $u(1) = 1$ and $u(0) = 0$, with 1 and 0 standing for the effective and the ineffective treatment, respectively). Verifiability is taken into account via a function $\psi(\cdot)$, by which vaccine A that is 80% effective in an i.i.d. manner is evaluated as $\psi(0.8)$, and vaccine B that is 80% effective for all is evaluated as $0.8\psi(1) + 0.2\psi(0)$. Notice that, A has to be viewed as a policy (and hence its single outcome's evaluation 0.8 needs to be transformed by $\psi(\cdot)$) to be comparable with B. The shape of $\psi(\cdot)$ captures the planner's attitude toward outcome uncertainty. By Jensen's inequality, a concave $\psi(\cdot)$, for example, favors vaccine A for its robustness (i.e., both vaccines present each citizen with 80% chances of effective vaccination ex ante, but A has only one outcome, or, put in another way, A is

[10]In this paper utilitarianism is a consequence of Proposition 1. Hence, unlike standard justifications that are by ethical principles or preferences aggregation, here it is justified by the idea of symmetry applied to policy outcomes of different lengths. And the result applies to qualitative prizes as well.

verifiable). Along the same line, a strong concavity of $\psi(\cdot)$ can potentially accommodate some planner's aversion to policies that may lead to catastrophic outcomes, even when they are unlikely (Fleurbaey 2018). Indeed, even if vaccine B becomes 99% chances effective for all, a strongly concave $\psi(\cdot)$ weighs highly the 1% scenario where the vaccine is not effective at all, and the planner may still prefer vaccine A despite the lower effective rate ex ante.

The monetary case is more nuanced because the utility function $u(\cdot)$, which was degenerate in the vaccine case, now captures the planner's attitude toward identity uncertainty. A concave $u(\cdot)$ reflects aversion to identity uncertainty and behaviorally it implies that the planner is more willing to sacrifice efficiency (i.e., average income) for equality. This can be seen in the utilitarian evaluation of an income profile $\mathbf{x} = (x_1, x_2, ..., x_I)$ by $U(\mathbf{x}) = \frac{1}{I}\sum_{i=1}^{I} u(x_i)$. A policy f as a state-contingent income profile is evaluated by $\sum_{s\in\mathcal{S}} \psi\left(U(f(s))\right) P(s)$, with P capturing the planner's subjective probability assessments of policy outcomes and function $\psi(\cdot)$, attitude toward outcome uncertainty. Under a concave $\psi(\cdot)$, the planner again dislikes outcome uncertainty and tends to value a policy's low variability of state-contingent utilitarian welfare.

3 Conclusion

We argue that real-world uncertainties are known when they are verifiable, as in the mechanics of a roulette wheel or the program of a computer simulation. Verifiable uncertainty is introduced in this light as a real-world representation of risk. We propose to formally model it by classical lottery — a finite collection of equally likely prizes. By substituting classical lottery for the risk part of the smooth model of ambiguity, we provide a model that assesses unverifiable general uncertainty by combining various scenarios of how it can be reduced to verifiable uncertainty, which we argue people often do in practice. The model allows for different attitudes across the two domains.

Our technical efforts seek to provide a normative foundation for this model. We first adapt the axioms of vNM expected utility to the domain of classical lottery, and obtain a counterpart of the model which we call average verifiable utility. Then, for any general prospect, we propose a set of desiderata the agent might want to follow in finding its indifferent classical lottery. These desiderata, mostly addressed to two-prize prospects, convey the normative content in combining various scenarios of "m out of n chances of winning a prize". Results are obtained for each prize pair and they are integrated together to pin down the general mode, which also includes the evaluation of non-two-prize prospects as a corollary. The setup we use here has the essential richness of the popular Anscombe-Aumann framework, and with our approach it can potentially host many other models.

References

Ahn, D. S. (2008): "Ambiguity without a state space," *The Review of Economic Studies*, 75, 3–28. 12

Allais, M. (1953): "L'extension des théories de l'équilibre économique général et du rendement social au cas du risque," *Econometrica, Journal of the Econometric Society*, 269–290. 4

Anscombe, F. J., and R. J. Aumann (1963): "A definition of subjective probability," *Annals of mathematical statistics*, 34, 199–205. 2

Debreu, G., and T. C. Koopmans (1982): "Additively decomposed quasiconvex functions," *Mathematical Programming*, 24, 1–38. 14, 29

Denti, T., and L. Pomatto (2022): "Model and predictive uncertainty: A foundation for smooth ambiguity preferences," *Econometrica*, 90, 551–584. 12

Ergin, H., and F. Gul (2009): "A theory of subjective compound lotteries," *Journal of Economic Theory*, 144, 899–929. 13

Flanigan, B., P. Gölz, A. Gupta, B. Hennig, and A. D. Procaccia (2021): "Fair algorithms for selecting citizens' assemblies," *Nature*, 596, 548–552. 3

Fleurbaey, M. (2018): "Welfare economics, risk and uncertainty," *Canadian Journal of Economics/Revue canadienne d'économique*, 51, 5–40. 16

Flores-Szwagrzak, K. (2024): "Learning by Convex Combination," working paper, Copenhagen Business School [wp], Denmark. 13

Gigerenzer, G. (1991): "How to make cognitive illusions disappear: Beyond heuristics and biases," *European review of social psychology*, 2, 83–115. 4

——— (1996): "On narrow norms and vague heuristics: A reply to Kahneman and Tversky.." 4

Gilboa, I., and D. Schmeidler (1989): "Maxmin expected utility with non-unique prior," *Journal of mathematical economics*, 18, 141–153. 12, 13

Grant, S., A. Kajii, B. Polak, and Z. Safra (2010): "Generalized utilitarianism and Harsanyi's impartial observer theorem," *Econometrica*, 78, 1939–1971. 15

Harsanyi, J. C. (1955): "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility," *Journal of political economy*, 63, 309–321. 3, 15

He, Y. (2021): "Revisiting Ellsberg's and Machina's paradoxes: a two-stage evaluation model under ambiguity," *Management Science*, 67, 6897–6914. 6, 13

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–91. 4

Keeney, R. L., and H. Raiffa (1993): *Decisions with multiple objectives: preferences and value trade-offs*: Cambridge university press. 11, 13

Keynes, J. M. (1937): "The general theory of employment," *The quarterly journal of economics*, 51, 209–223. 2

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73, 1849–1892. 2, 12

Knight, F. H. (1921): *Risk, uncertainty and profit* Volume 31: Houghton Mifflin. 2

Kothiyal, A., V. Spinu, and P. P. Wakker (2014): "Average utility maximization: a preference foundation," *Operations Research*, 62, 207–218. 7, 8

Kreps, D. M., and E. L. Porteus (1978): "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica: journal of the Econometric Society*, 185–200. 12

Marquis de Laplace, P. S. (1902): *A philosophical essay on probabilities*: Wiley. 1

Letsou, C., S. Naeh, and U. Segal (2022): "All probabilities are equal, but some probabilities are more equal than others," *Economic Theory*, 74, 423–445. 3

Li, C., K. Rohde, and P. P. Wakker (2023): "The deceptive beauty of monotonicity, and the million-dollar question: row-first or column-first aggregation," *Working Paper, Erasmus University Rotterdam*. 6

Lipkus, I. M., and J. G. Hollands (1999): "The visual communication of risk," *JNCI Monographs*, 1999, 149–163. 4

Marinacci, M. (2015): "Model uncertainty," *Journal of the European Economic Association*, 13, 1022–1100. 5

Nau, R. F. (2006): "Uncertainty aversion with second-order utilities and probabilities," *Management science*, 52, 136–145. 2, 13

von Neumann, J., and O. Morgenstern (1944): *Theory of Games and Economic Behavior*: Princeton University Press. 7

Procaccia, A. D. (2022): "A More Perfect Algorithm," *Scientific American*, 327, 52–59. 3

Quiggin, J. (1982): "A theory of anticipated utility," *Journal of Economic Behavior & Organization*, 3, 323–343. 13

Rawls, J. (2001): *Justice as fairness: A restatement*: Harvard University Press. 3, 15

Savage, L. J. (1954): *The foundations of statistics*: New York: John Wiley and Sons. (Second addition in 1972, Dover). 5, 11

Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–87. 12, 13

Segal, U. (1987): "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28, 175–202. 12

——— (1990): "Two-stage lotteries without the reduction axiom," *Econometrica: Journal of the Econometric Society*, 349–377. 12

Seo, K. (2009): "Ambiguity and second-order belief," *Econometrica*, 77, 1575–1605. 12

Smith, J. E., and J. S. Dyer (2021): "On (measurable) multiattribute value functions: An expository argument," *Decision Analysis*, 18, 247–256. 11

Spiegelhalter, D., M. Pearson, and I. Short (2011): "Visualizing uncertainty about the future," *Science*, 333, 1393–1400. 3, 4

Wakker, P. P. (1989): *Additive Representations of Preferences: A New Foundation of Decision Analysis*: Dordorecht. 10, 12, 25

Wang, F. (2019): "Comparative Ambiguity Attitudes," *working paper*. 8

——— (2022): "Rank-Dependent Utility Under Multiple Priors," *Management Science*, 68, 8166–8183. 13

4 Appendix

4.1 Proof of Proposition 1

The proof consists of three steps. In Step 1, we establish certain properties that allow us to identify classical lottery as probability distribution with rational probability numbers. Correspondingly, the preference over classical lotteries carries over to a preference over rational probability distributions. Step 2 shows that the new preference satisfies von Neumann-Morgenstern axioms (restricted on the domain of rational probability distributions). Finally, in Step 3, we establish the expected utility model for the new preference, which can be translated back to the classical lotteries domain and that delivers the average verifiable utility model. The lack of a continuum of probability numbers presents the main challenge. Details are below.

Proof. Step 1: Axioms 1–3 imply two intuitive properties that allow us to identify classical lottery as rational-numbered probability distribution.

Anonymity: For all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, if $\mathbf{x}$ is a permutation of $\mathbf{y}$ (that is, $|\mathbf{x}| = |\mathbf{y}|$ and there is a bijection ρ on $\{1, 2, ..., |\mathbf{x}|\}$ such that $\mathbf{x}(i) = \mathbf{y}(\rho(i))$ for all $1 \leq i \leq |\mathbf{x}|$), then $\mathbf{x} \sim \mathbf{y}$. We shall prove by induction. If $|\mathbf{x}| = 1$, we have $\mathbf{x} = \mathbf{y}$ and hence $\mathbf{x} \sim \mathbf{y}$ trivially. Suppose that the claim holds for all $\mathbf{x}$ and $\mathbf{y}$ with $|\mathbf{x}| = |\mathbf{y}| \leq n$. Then for $|\mathbf{x}| = |\mathbf{y}| = n+1$, pick any length-$n$ component of $\mathbf{x}$, which is indifferent to a corresponding length-n component of $\mathbf{y}$, and the remaining prize in $\mathbf{x}$ should coincide with that in $\mathbf{y}$, and hence the Cancellation axiom is applicable to deliver $\mathbf{x} \sim \mathbf{y}$.

Replicability: For all $\mathbf{x} \in \mathcal{X}$, $\mathbf{x} \sim \mathbf{x}^n$ for all $n \geq 1$. If $\mathbf{x} \succ \mathbf{x}^n$ for some $\mathbf{x}$ and n, then Cancellation implies $\mathbf{x}^k \succ \mathbf{x}^{k \cdot n}$ for all $k \geq 1$. In particular, by taking $k = n^l$ with $l \geq 1$, we have $\mathbf{x} \succ \mathbf{x}^n \succ \mathbf{x}^{n^2} \succ \cdots \succ \mathbf{x}^{n^l} \succ \ldots$. With $\mathbf{x} \succ \mathbf{x}^n$, we check Archimedeanity by mixing $\mathbf{x}$ with $\mathbf{x}$ itself. But, for any integer $m \geq 1$, consider $h = n^l - 1$ for any l large enough so that $h \geq m$, Archimedeanity is violated because we have $\mathbf{x}^n \succ \mathbf{x}^{n^l} = \mathbf{x}^h \vee \mathbf{x}$. If instead $\mathbf{x}^n \succ \mathbf{x}$ for some $\mathbf{x}$ and n, similar argument leads to $\cdots \succ \mathbf{x}^{n^l} \succ \cdots \succ \mathbf{x}^{n^2} \succ \mathbf{x}^n \succ \mathbf{x}$, which also violates Archimedeanity as we mix $\mathbf{x}$ with itself for $\mathbf{x}^n \succ \mathbf{x}$.

Let $\Delta_0(X)$ denote the set of probability distributions on X with finite support and let $\Delta_{0,\mathbb{Q}}(X) = \{p \in \Delta_0(X) : p(x) \in \mathbb{Q}_{[0,1]}, \forall x \in X\}$ be the subset of $\Delta_0(X)$ that only take rational numbers as probabilities. Here we adopt the notation that $\mathbb{Q}_I$ denotes the set of rational numbers in set $I \subset \mathbb{R}$. A generic $\mathbf{p} \in \Delta_{0,\mathbb{Q}}(X)$ is in the form of $(x_1, p_1; ...; x_k, p_N)$ for some $\{x_n\}_{n=1}^N \subset X$ and some $\{p_n\}_{n=1}^N \subset \mathbb{Q}_{(0,1]}$ with $\sum_{n=1}^N p_n = 1$, and the interpretation is that prize x_n happens with probability $\mathbf{p}(x_n) = p_n$ for all n. For any classical lottery $\mathbf{x}$, let $\mathbf{p_x} \in \Delta_{0,\mathbb{Q}}(X)$ denote the frequencies of the prizes in $\mathbf{x}$, and, for any $\mathbf{p} \in \Delta_{0,\mathbb{Q}}(X)$, let $H(\mathbf{p}) = \{\mathbf{x} \in \mathcal{X} : \mathbf{p_x} = \mathbf{p}\}$ be the set of classical lotteries whose frequencies of prizes coincide with $\mathbf{p}$.

The preference $\succcurlyeq$ on $\mathcal{X}$ can be transplanted to $\Delta_{0,\mathbb{Q}}(X)$ and form a new binary relation $\succcurlyeq^*$. For this we first show that for all $\mathbf{p} \in \Delta_{0,\mathbb{Q}}(X)$ and all $\mathbf{x}, \mathbf{y} \in H(\mathbf{p})$, $\mathbf{x} \sim \mathbf{y}$. To see why, let the support of p be denoted by $\{x_1, ..., x_k\}$ and, for each $i \in \{1, 2, ..., k\}$, let $\mathbf{p}(x_i)$ be denoted by $\frac{m_i}{n_i}$ such that the greatest common divisor of m_i and n_i is 1. Let N be the least common multiple of $\{n_1, n_2, ..., n_k\}$. Then the classical lottery $\mathbf{x}^o = (x_1, ..., x_1, x_2, ...x_2, ..., x_k, ..., x_k)$ consisting of $\frac{N}{n_i} m_i$ many x_i with $i \in \{1, 2, ..., k\}$ must be in $H(\mathbf{p})$, and, any other $\mathbf{x} \in H(\mathbf{p})$ must be a permutation of a finite concatenation of $\mathbf{x}^o$. Anonymity and replicability imply that they are all indifferent to each other. We can therefore define that for all $\mathbf{p}, \mathbf{q} \in \Delta_{0,\mathbb{Q}}(X)$, $\mathbf{p} \succcurlyeq^* \mathbf{q}$ if and only if $\mathbf{x} \succcurlyeq \mathbf{y}$ for some $\mathbf{x} \in H(\mathbf{p})$ and $\mathbf{y} \in H(\mathbf{q})$. The resulting binary relation $\succcurlyeq^*$ on $\Delta_{0,\mathbb{Q}}(X)$ is complete and transitive because $\succcurlyeq$ is. The symmetric and asymmetric parts of $\succcurlyeq^*$ will be similarly denoted by $\succ^*$ and $\sim^*$, respectively.

Step 2: The preference $\succcurlyeq^*$ on $\Delta_{0,\mathbb{Q}}(X)$ satisfies von Neumann-Morgenstern axioms.

For vNM independence, we want to show, for all $\mathbf{p}, \mathbf{p}', \mathbf{q} \in \Delta_{0,\mathbb{Q}}(X)$ and rational $r \in$

$\mathbb{Q}_{(0,1)}$, $\mathbf{p} \succcurlyeq^* \mathbf{p}'$ if and only if $(1-r)\mathbf{p} + r\mathbf{q} \succcurlyeq^* (1-r)\mathbf{p}' + r\mathbf{q}$. Here we restrict to rational r so that $(1-r)\mathbf{p} + r\mathbf{q}$ is in $\Delta_{0,\mathbb{Q}}(X)$. Notice that we can always pick $\mathbf{x} \in H(\mathbf{p})$, $\mathbf{x}' \in H(\mathbf{p}')$, and $\mathbf{y} \in H(\mathbf{q})$ such that they all belong to the same X^k for some k, because we can take any $\mathbf{x}_0 \in H(\mathbf{p})$, $\mathbf{x}'_0 \in H(\mathbf{p}')$, and $\mathbf{y}_0 \in H(\mathbf{q})$ with lengths k_1, k_2, and k_3, respectively, and take $\mathbf{x}$, $\mathbf{x}'$, and $\mathbf{y}$ to be $\mathbf{x}_0^{k_2k_3}$, $(\mathbf{x}'_0)^{k_1k_3}$, and $\mathbf{y}_0^{k_1k_2}$, respectively, with $k = k_1k_2k_3$. For any rational number $r = \frac{m}{n}$ in $(0,1)$, by Anonymity and Cancellation we have $\mathbf{x} \succcurlyeq \mathbf{x}'$ if and only if $(\mathbf{x}^{n-m} \vee \mathbf{y}^m) \succcurlyeq ((\mathbf{x}')^{n-m} \vee \mathbf{y}^m)$. Since $\mathbf{x}^{n-m} \vee \mathbf{y}^m \in H((1-r)\mathbf{p} + r\mathbf{q})$ and $(\mathbf{x}')^{n-m} \vee \mathbf{y}^m \in H((1-r)\mathbf{p}' + r\mathbf{q})$, the axiom has to hold.

As to vNM continuity, we want to show, for any $\mathbf{p}, \mathbf{p}', \mathbf{q} \in \Delta_{0,\mathbb{Q}}(X)$ with $\mathbf{p} \succ^* \mathbf{p}'$, the existence of rational $r, r' \in \mathbb{Q}_{(0,1)}$ such that $(1-r)\mathbf{p} + r\mathbf{q} \succ^* \mathbf{p}'$ and $\mathbf{p} \succ^* (1-r')\mathbf{p}' + r'\mathbf{q}$. Pick $\mathbf{x} \in H(\mathbf{p})$, $\mathbf{x}' \in H(\mathbf{p}')$, and $\mathbf{y} \in H(\mathbf{q})$. By Archimedeanity, there exists an $N \in \mathbb{N}$, such that $(\mathbf{x}^n \vee \mathbf{y}) \succ \mathbf{x}'$ and $\mathbf{x} \succ ((\mathbf{x}')^n \vee \mathbf{y})$ for all $n \geq N$. This means that for $r = \frac{1}{n+1}$ with $n \geq N$ we have $(1-r)\mathbf{p} + r\mathbf{q} \succ^* \mathbf{p}'$ and $\mathbf{p} \succ^* (1-r)\mathbf{p}' + r\mathbf{q}$, and hence the axiom.

Step 3: Preference $\succcurlyeq^*$ on $\Delta_{0,\mathbb{Q}}(X)$ admits an expected utility representation, which corresponds to the average verifiable utility model for $\succcurlyeq$ on $\mathcal{X}$. The strategy of the proof is standard, by consecutively simplifying a general lottery into a binary lottery, though some technical complications are called for because we only have rational probabilities.

To start, for any two prizes $x, y \in X$ and any lottery $\mathbf{p} \in \Delta_{0,\mathbb{Q}}(X)$ with $x \succcurlyeq^* \mathbf{p} \succcurlyeq^* y$, let $\mathbf{p}^+(x,y) = \{r \in \mathbb{Q}_{[0,1]} : (x, r; y, 1-r) \succcurlyeq^* \mathbf{p}\}$ and $\mathbf{p}^-(x,y) = \{r \in \mathbb{Q}_{[0,1]} : \mathbf{p} \succcurlyeq^* (x, r; y, 1-r)\}$. By completeness, $\mathbf{p}^+(x,y) \cup \mathbf{p}^-(x,y) = \mathbb{Q}_{[0,1]}$, and, by independence and transitivity, we have $inf\,\mathbf{p}^+(x,y) = sup\,\mathbf{p}^-(x,y)$. We denote this common number by $m_{x,y}(\mathbf{p})$. Notice that, when $m_{x,y}(\mathbf{p})$ is rational, by continuity we must have $\mathbf{p} \sim^* (x, m_{x,y}(\mathbf{p}); y, 1-m_{x,y}(\mathbf{p}))$.

Fix any two prizes $x_1, x_0 \in X$ with $x_1 \succ^* x_0$, we claim that $\succcurlyeq^*$ restricted on $\Delta_{0,\mathbb{Q}}(\{x \in X : x_1 \succcurlyeq^* x \succcurlyeq^* x_0\})$ admits an expected utility representation. Define a utility function u on $\{x \in X : x_1 \succcurlyeq^* x \succcurlyeq^* x_0\}$ as follows. Let $u(x_1) = 1$ and $u(x_0) = 0$, and for any other $x \in X$ with $x_1 \succcurlyeq^* x \succcurlyeq^* x_0$, let $u(x) = m_{x_1,x_0}(x)$, where prize x is seen as a degenerate lottery. For any $\mathbf{p} \in \Delta_{0,\mathbb{Q}}(\{x \in X : x_1 \succcurlyeq^* x \succcurlyeq^* x_0\})$, by independence we can approximate every prize in the support of $\mathbf{p}$ by (binary and rational) probability distributions on $\{x_1, x_0\}$, from which we get

$$E_{\mathbf{p}}u = \sum m_{x_1,x_0}(x)\mathbf{p}(x) = m_{x_1,x_0}(\mathbf{p}).$$

To show this expected utility does represent $\succcurlyeq^*$, take any $\mathbf{p}, \mathbf{q} \in \Delta_{0,\mathbb{Q}}(\{x \in X : x_1 \succcurlyeq^* x \succcurlyeq^* x_0\})$ and start with the direction of $E_{\mathbf{p}}u > E_{\mathbf{q}}u$ implying $\mathbf{p} \succ^* \mathbf{q}$. This holds, because $m_{x_1,x_0}(\mathbf{p}) > m_{x_1,x_0}(\mathbf{q})$ implies the existence of $\mathbf{r}_1, \mathbf{r}_2 \in \Delta_{0,\mathbb{Q}}(\{x_1, x_0\})$ with $\mathbf{p} \succ^* \mathbf{r}_1 \succ^* \mathbf{r}_2 \succ^* \mathbf{q}$. For the direction of $\mathbf{p} \succ^* \mathbf{q}$ implying $E_{\mathbf{p}}u > E_{\mathbf{q}}u$, it suffices to show $m_{x_1,x_0}(\mathbf{p}) > m_{x_1,x_0}(\mathbf{q})$. Because $x_1 \succcurlyeq^* \mathbf{p} \succ^* \mathbf{q}$, there exists a rational $\alpha \in (0,1)$ such that $\mathbf{p}_1 \succ^* \alpha x_1 + (1-\alpha)\mathbf{q} \succ^* \mathbf{q}$ by continuity and independence. And, by independence again, this leads to $m_{x_1,x_0}(\mathbf{p}) \geq \alpha + (1-\alpha)m_{x_1,x_0}(\mathbf{q}) > m_{x_1,x_0}(\mathbf{q})$, where the second inequality follows from $m_{x_1,x_0}(\mathbf{q}) < 1$. Expected utility representation is hence established, with the utility function u being unique up to affine transformation.

For any other prizes x, x' with $x \succcurlyeq^* x_1 \succ^* x_0 \succcurlyeq^* x'$, we can repeat the arguments above to establish an expected utility model on $\Delta_{0,\mathbb{Q}}(\{x \in X : x \succcurlyeq^* x \succcurlyeq^* x'\})$ with utility function $u_{x,x'}$ on $\{x \in X : x \succcurlyeq^* x \succcurlyeq^* x'\}$. Because $\{x \in X : x_1 \succcurlyeq^* x \succcurlyeq^* x_0\} \subset \{x \in X : x \succcurlyeq^* x \succcurlyeq^* x'\}$ and by the uniqueness property of utility function u, it must be that, after some affine transformation, $u_{x,x'}$ restricted on $\{x \in X : x_1 \succcurlyeq^* x \succcurlyeq^* x_0\}$ coincides with u. Take this particular affine transformation of $u_{x,x'}$ and extend u to $\{x \in X : x \succcurlyeq^* x \succcurlyeq^* x'\}$ accordingly. We can extend u to the whole set X by considering all such pairs $x \succ^* x'$, because any two such extensions have to agree on their common part. Finally, we claim that expected utility model with this fully extended u represents $\succcurlyeq^*$ on $\Delta_{0,\mathbb{Q}}(X)$. This is true because, for any $\mathbf{p}, \mathbf{q} \in \Delta_{0,\mathbb{Q}}(X)$, the model represents $\succcurlyeq^*$ restricted on $\Delta_{0,\mathbb{Q}}(\{x \in X : \overline{x} \succcurlyeq^* x \succcurlyeq^* \underline{x}\})$, where $\overline{x}$ and $\underline{x}$ are (possibly one of) the best and the worst prize in $supp(\mathbf{p}) \cup supp(\mathbf{q}) \cup \{x_1, x_0\}$, respectively. □

4.2 Proof of Lemma 3

Let us generalize the notion of matching frequency to general acts. For any two prizes $x \succ y$ and any act $f \in \mathcal{F}$ with $x \succcurlyeq f \succcurlyeq y$, the matching frequency of f is

$$m_{x,y}(f) = sup\{r_x(\mathbf{x}) : f \succcurlyeq \mathbf{x}, \mathbf{x} \in \mathcal{X}_{x,y}\}. \tag{4.1}$$

We shall prove the following result, which is a version of Lemma 3 that is also applicable to general acts.

Lemma 7. *Under Axioms 1–3 and 5, for any two prizes $x \succ y$ and any two acts $f, g \in \mathcal{F}$ with $x \succcurlyeq f \succcurlyeq y$ and $x \succcurlyeq g \succcurlyeq y$,*

$$f \succcurlyeq g \Longleftrightarrow m_{x,y}(f) \geq m_{x,y}(g).$$

Proof. By applying Denseness twice we can conclude that for any two acts $f \succ g$ there exist classical lotteries $\mathbf{x}, \mathbf{y}$ such that $f \succ \mathbf{x} \succ \mathbf{y} \succ g$.

Consider any two prizes $x \succ y$ and any two acts f, g with $x \succcurlyeq f, g \succcurlyeq y$. If $f \sim g$, obviously we have $m_{x,y}(f) = m_{x,y}(g)$. Now suppose that $f \succ g$. By the above claim there exist classical lotteries $\mathbf{x}, \mathbf{y}$ such that $f \succ \mathbf{x} \succ \mathbf{y} \succ g$, and, by the average utility representation, there exists two-prize classical lotteries $\mathbf{z_1}, \mathbf{z_2} \in \mathcal{X}_{x,y}$ such that $\mathbf{x} \succ \mathbf{z_1} \succ \mathbf{z_2} \succ \mathbf{y}$, where $r_x(\mathbf{z_1}) = q_1 > r_x(\mathbf{z_2}) = q_2$. Hence we have $m_{x,y}(f) > m_{x,y}(g)$ (indeed, $m_{x,y}(f) > q_1 > q_2 > m_{x,y}(g)$). □

Axioms 1–4 imply that for all prizes $x \succ y$ and any act $f \in \mathcal{F}_{x,y}$, we have $x \succcurlyeq f \succcurlyeq y$. Hence, under Axioms 1–5, Lemma 3 can be seen as a special case of Lemma 7.

We finish this section by discussing an example that satisfies Axioms 1–4 but not Denseness. Suppose $\mathcal{S} = \{s_1, s_2\}$, $X = \{\$0, \$100\}$, and $\succsim$ on $\mathcal{F}$ (which contains only binary acts) is lexicographic: $f \succsim g$ if and only if $f(s_1)$ has a higher frequency of \$100 than $g(s_1)$ or, in the case they are equal, $f(s_2)$ has a higher frequency of \$100 than $g(s_2)$. Here the matching frequency of any act $f \in \mathcal{F}$ is equal to the frequency $f(s_1)$ has for \$100. Then, in the case $f \succ \mathbf{x}$ (or $\mathbf{x} \succ f$) with classical lottery $\mathbf{x}$ being f's matching classical lottery, we cannot find another classical lottery $\mathbf{y}$ such that $f \succ \mathbf{y} \succ \mathbf{x}$ (respectively, $\mathbf{x} \succ \mathbf{y} \succ f$). Notice that here the matching frequency is continuous on $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$, that is, Axiom 6 is also satisfied.

4.3 Proof of Theorem 4

Before proving the theorem, we would like to briefly comment on why Axiom 7 is crucial. When applied to a fixed state s, the axiom requires that the relative importances of frequency increments to be independent of what is happening outside state s, conveying an idea of separability. When the axiom is applied across states, it demands the frequencies on different states to be valued in the same scale, suggesting a flavor of cardinal equivalence. However, different weights are still allowed to be attached to different states.

The sufficiency proof can be divided into two parts, with the first part establishing the model for any fixed prize pair, and the second part integrating the models on different prize pairs into a unified one. More specifically, for any prize pair $x \succ y$, Step 1.1 identifies preference $\succcurlyeq$ over two-prize acts $\mathcal{F}_{x,y}$ as a preference $\succcurlyeq_{x,y}$ on $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$, the set of state-contingent frequencies of winning x. Step 1.2 shows that $\succcurlyeq_{x,y}$ is continuous and has a unique continuous extension to $[0,1]^{\mathcal{S}}$. Step 1.3 shows that after extension the consistent reduction axiom still holds on $(0,1)^{\mathcal{S}}$. This turns out to be sufficient for additive representation $V_{x,y}$ on $[0,1]^{\mathcal{S}}$, as shown in Step 1.4, and, as a consequence, the target model for $\mathcal{F}_{x,y}$. Step 1.5 shows that, for general acts that can be bounded in preference terms by acts in $\mathcal{F}_{x,y}$, the preference over them (now identified as vectors in $[0,1]^{\mathcal{S}}$) can also be represented by $V_{x,y}$. This concludes the first part. In the second part, Step 2.1 shows that, for any four prizes $x' \succcurlyeq x \succ y \succcurlyeq y'$, the model $V_{x',y'}$ is cardinally equivalent to model $V_{x,y}$ for the evaluations of acts in $\mathcal{F}_{x,y}$ (as $\mathcal{F}_{x',y'}$ bounds $\mathcal{F}_{x,y}$, model $V_{x',y'}$ can be applied to evaluate $\mathcal{F}_{x,y}$). This allows us to glue the models $\{V_{x,y}\}_{x,y\in\mathcal{X}}$ together for a unified model V, as in Step 2.2, by enlarging the preference gap $x \succ y$ to cover all prizes. In Step 2.3, the model V is then translated back to the domain $\mathcal{F}$ to deliver the target model.

Proof. We shall focus on the sufficiency of the axioms, as their necessity is straightforward. And, if $x \sim y$ for all prizes x and y, Axioms 1–4 imply that $f \sim g$ for all acts f and g. In this degenerate case, the model holds along with the uniqueness properties. We henceforth assumes that $\succcurlyeq$ is not total indifference over prizes.

Step 1.1: For any two prizes $x \succ y$, preference $\succcurlyeq$ over $\mathcal{F}_{x,y}$ can be seen as a weak-order preference $\succcurlyeq_{x,y}$ on $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$. For any (rational) vectors $\mathbf{q}_1, \mathbf{q}_2 \in \mathbb{Q}^{\mathcal{S}}_{[0,1]}$, we say $\mathbf{q}_1$ is preferred to $\mathbf{q}_2$ on prize-pair (x, y), denoted by $\mathbf{q}_1 \succcurlyeq_{x,y} \mathbf{q}_2$, if there exist acts $f_1, f_2 \in \mathcal{F}_{x,y}$ such that $\mathbf{q}_1(s) = r_x(f_1(s))$ and $\mathbf{q}_2(s) = r_x(f_2(s))$ for all $s \in \mathcal{S}$ and $f_1 \succcurlyeq f_2$. Under Axioms 1–4, the representing binary acts of any $\mathbf{q} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]}$ are indifferent, and we have $\mathbf{q}_1 \succcurlyeq_{x,y} \mathbf{q}_2$ if and only if $f_1 \succcurlyeq f_2$ for *all* $f_1, f_2 \in \mathcal{F}_{x,y}$ that represent $\mathbf{q}_1$ and $\mathbf{q}_2$, respectively. Moreover, preference $\succsim_{x,y}$ is complete and transitive. The symmetric and asymmetric parts of $\succsim_{x,y}$ are denoted by $\sim_{x,y}$ and $\succ_{x,y}$, respectively. All the axioms on $\succsim$ are inherited by $\succsim_{x,y}$. And recall that the matching frequency of any $\mathbf{q} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]}$, denoted by $m_{xy}(\mathbf{q})$, refers to the matching frequency $m_{x,y}(f)$ of any binary act $f \in$ that $\mathbf{q}$ represents.

Step 1.2: We claim that $\succsim_{x,y}$ is continuous in the sense that for all $\mathbf{q} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]}$ the sets $\{\mathbf{r} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]} : \mathbf{r} \succ_{x,y} \mathbf{q}\}$ and $\{\mathbf{r} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]} : \mathbf{q} \succ_{x,y} \mathbf{r}\}$ are open in $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$ under the usual metric topology. Take any $\mathbf{q}' \succ_{x,y} \mathbf{q}$, by Continuity we can mix $\mathbf{q}'$ with the worst vector $\mathbf{0}$ and then there should exist an $\alpha \in \mathbb{Q}_{[0,1]}$ with $\alpha\mathbf{q}' + (1-\alpha)\mathbf{0} \succ_{x,y} \mathbf{q}$. Let $N_{\mathbf{q}'} \subset \mathcal{S}$ denote the set of states on which $\mathbf{q}'$ assigns 0. By monotonicity, every member of $\{\mathbf{r} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]} : \mathbf{r} > \alpha\mathbf{q}' + (1-\alpha)\mathbf{0} \text{ } on \text{ } S \backslash N_{\mathbf{q}'}, \mathbf{r} \geq 0 \text{ } on \text{ } N_{\mathbf{q}'}\}$ is strictly preferred to $\mathbf{q}$, and moreover, the set is open and contains $\mathbf{q}'$. Since $\mathbf{q}'$ is arbitrary, set $\{\mathbf{r} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]} : \mathbf{r} \succ_{x,y} \mathbf{q}\}$ is open. Similarly, by mixing $\mathbf{q}$ with the best vector $\mathbf{1}$, set $\{\mathbf{r} \in \mathbb{Q}^{\mathcal{S}}_{[0,1]} : \mathbf{q} \succ_{x,y} \mathbf{r}\}$ can also be shown to be open.

We can (continuously) extend $\succsim_{x,y}$ from $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$ to $[0,1]^{\mathcal{S}}$ by monotonicity. For this, we claim that for any $\mathbf{q} \in [0,1]^{\mathcal{S}}$ there exists an $m_{\mathbf{q}} \in [0,1]$ such that $lim_{n\to\infty} m_{x,y}(\mathbf{q}_n) = m_{\mathbf{q}}$ for all sequence $\mathbf{q}_n \to \mathbf{q}$. If this is not the case and there are two sequences $\mathbf{r}_n \to \mathbf{q}$ and $\mathbf{r}'_n \to \mathbf{q}$ but $lim_{n\to\infty} m_{x,y}(\mathbf{r}_n) > lim_{n\to\infty} m_{x,y}(\mathbf{r}'_n)$. The sequence $\{\hat{\mathbf{r}}_n\}$ with $\hat{\mathbf{r}}_n = \mathbf{r}_{n/2}$ when n is even and $\hat{\mathbf{r}}_n = \mathbf{r}_{(n+1)/2}$ when n is odd represents a convergent sequence of acts on $\mathcal{F}_{x,y}$ and that would violate the continuity axiom. Hence the claim is valid, and, we can therefore define the matching frequency of $\mathbf{q}$ as $m_{x,y}(\mathbf{q}) = m_{\mathbf{q}}$. The extended preference $\bar{\succsim}_{x,y}$ is defined following the natural order on matching frequencies, with $\bar{\succ}_{x,y}$ being the strict part of $\bar{\succsim}_{x,y}$. Note that $\bar{\succsim}_{x,y}$ is continuous on $[0,1]^{\mathcal{S}}$, and it respects strict monotonicity on essential states and total indifference on inessential states.

Step 1.3: We claim that the Consistent Reduction axiom is preserved for the extended preference $\bar{\succsim}_{x,y}$ restricted on $(0,1)^{\mathcal{S}}$. Suppose that, on prize pair $x \succ y$, some preference pair shows the incremental from a to $a+\delta$ weighs larger than that from a' to $a'+\delta'$, while another preference pair shows the incremental from a' to $a'+\delta'$ weighs strictly larger than that from a to $a+\delta$. Here $a, a+\delta, a', a'+\delta' \in (0,1)$ are real numbers, not necessarily rational. Take $b = a+\delta$ and $b' = a'+\delta'$, and, by definition there exist $\mathbf{p}, \mathbf{q}, \mathbf{p}', \mathbf{q}' \in (0,1)^{\mathcal{S}}$ and essential states s and s', such that $m_{x,y}(\mathbf{a}_s\mathbf{p}) \geq m_{x,y}(\mathbf{b}_s\mathbf{q})$, $m_{x,y}(\mathbf{a}'_s\mathbf{p}) \leq m_{x,y}(\mathbf{b}'_s\mathbf{q})$, $m_{x,y}(\mathbf{a}_{s'}\mathbf{p}') \leq m_{x,y}(\mathbf{b}_{s'}\mathbf{q}')$, but $m_{x,y}(\mathbf{a}'_{s'}\mathbf{p}') > m_{x,y}(\mathbf{b}'_{s'}\mathbf{q}')$. We shall exploit monotonicity and continuity to find a corresponding violation for $\succsim_{x,y}$ on $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$. Starting with $m_{x,y}(\mathbf{a}'_{s'}\mathbf{p}') > m_{x,y}(\mathbf{b}'_{s'}\mathbf{q}')$, we can find rational $\mathbf{a}'^{,\mathbb{Q}} < \mathbf{a}'$, $\mathbf{p}'^{,\mathbb{Q}} < \mathbf{p}'$, $\mathbf{b}'^{,\mathbb{Q}} > \mathbf{b}'$, and $\mathbf{q}'^{,\mathbb{Q}} > \mathbf{q}'$ such that $m_{x,y}(\mathbf{a}'^{,\mathbb{Q}}_{s'}\mathbf{p}'^{,\mathbb{Q}}) > m_{x,y}(\mathbf{b}'^{,\mathbb{Q}}_{s'}\mathbf{q}'^{,\mathbb{Q}})$. By $m_{x,y}(\mathbf{a}'_s\mathbf{p}) \leq m_{x,y}(\mathbf{b}'_s\mathbf{q})$ and s being essential, we have $m_{x,y}(\mathbf{a}'^{,\mathbb{Q}}_s\mathbf{p}) < m_{x,y}(\mathbf{b}'^{,\mathbb{Q}}_s\mathbf{q})$ and therefore we can find rational $\mathbf{p}^{\mathbb{Q}} > \mathbf{p}$ and $\mathbf{q}^{\mathbb{Q}} < \mathbf{q}$ such that $m_{x,y}(\mathbf{a}'^{,\mathbb{Q}}_s\mathbf{p}^{\mathbb{Q}}) < m_{x,y}(\mathbf{b}'^{,\mathbb{Q}}_s\mathbf{q}^{\mathbb{Q}})$. Notice that either event $\{s\}^c$ or event $\{s'\}^c$ has to be essential (i.e., contain at least one essential state) because we would otherwise have $\mathbf{a}' \leq \mathbf{b}'$ and $\mathbf{a}' > \mathbf{b}'$. Indeed, both event $\{s\}^c$ and event $\{s'\}^c$ have to be essential because, if $s = s'$, then s^c and $(s')^c$ are the same event, and, if $s \neq s'$, they have to be essential by definition. Since $\{s\}^c$ is essential, we will have $m_{x,y}(\mathbf{a}_s\mathbf{p}^{\mathbb{Q}}) > m_{x,y}(\mathbf{b}_s\mathbf{q}^{\mathbb{Q}})$ and we can therefore find rational $\mathbf{a}^{\mathbb{Q}} < \mathbf{a}, \mathbf{b}^{\mathbb{Q}} > \mathbf{b}$ that satisfy both $m_{x,y}(\mathbf{a}^{\mathbb{Q}}_s\mathbf{p}^{\mathbb{Q}}) > m_{x,y}(\mathbf{b}^{\mathbb{Q}}_s\mathbf{q}^{\mathbb{Q}})$ and $m_{x,y}(\mathbf{a}^{\mathbb{Q}}_{s'}\mathbf{p}'^{\mathbb{Q}}) < m_{x,y}(\mathbf{b}^{\mathbb{Q}}_{s'}\mathbf{q}'^{\mathbb{Q}})$.[11] Hence we end up with a violation of Consistent Reduction for $\succcurlyeq_{x,y}$ on $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$.

Step 1.4: The extended preference $\bar{\succcurlyeq}_{x,y}$ restricted on $(0,1)^{\mathcal{S}}$ can be represented by an additive model (Wakker 1989, Theorem IV 2.7), that is, for all $\mathbf{p}, \mathbf{q} \in (0,1)^{\mathcal{S}}$

$$\mathbf{p}\bar{\succcurlyeq}_{x,y}\mathbf{q} \Leftrightarrow \int_{\mathcal{S}} V_{x,y}(\mathbf{p}(s))dP_{x,y} \geq \int_{\mathcal{S}} V_{x,y}(\mathbf{q}(s))dP_{x,y}, \tag{4.2}$$

where $V_{x,y} : (0,1) \to \mathbb{R}$ is a continuous, strictly increasing function and $P_{x,y}$ is a probability measure on $\mathcal{S}$. The function $V_{x,y}$ is unique up to affine transformation and $P_{x,y}$ is unique. By monotonicity $V_{x,y}$ admits a unique continuous extension to $[0,1]$, which shall still be denoted by $V_{x,y}$.

We claim that model 4.2 with the extended $V_{x,y}$ actually represents $\bar{\succcurlyeq}_{x,y}$ on the full domain, $[0,1]^{\mathcal{S}}$. For all $\mathbf{p}, \mathbf{q} \in [0,1]^{\mathcal{S}}$ with $\mathbf{p}\bar{\succ}_{x,y}\mathbf{q}$, pick two convergent rational sequences in $(0,1)^{\mathcal{S}}$ with $\mathbf{p}_n$ converging to $\mathbf{p}$ and $\mathbf{q}_n$ converging to $\mathbf{q}$. By the definition of $\bar{\succcurlyeq}_{x,y}$, we have $lim_{n\to\infty} m_{x,y}(\mathbf{p}_n) = m_{x,y}(\mathbf{p}) > m_{x,y}(\mathbf{q}) = lim_{n\to\infty} m_{x,y}(\mathbf{q}_n)$. And, because model 4.2 represents $\bar{\succcurlyeq}_{x,y}$ restricted on $(0,1)^{\mathcal{S}}$, we have that, for all $a, b \in (lim_{n\to\infty} m_{x,y}(\mathbf{q}_n), lim_{n\to\infty} m_{x,y}(\mathbf{p}_n))$ with $a > b$,

$$\int_{\mathcal{S}} V_{x,y}(\mathbf{p}_n(s))dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(a)dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(b)dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(\mathbf{q}_n(s))dP_{x,y},$$

for all large n. By continuity, we have $\int_{\mathcal{S}} V_{x,y}(\mathbf{p}(s))dP_{x,y} = lim_{n\to\infty} \int_{\mathcal{S}} V_{x,y}(\mathbf{p}_n(s))dP_{x,y}$ and $\int_{\mathcal{S}} V_{x,y}(\mathbf{q}(s))dP_{x,y} = lim_{n\to\infty} \int_{\mathcal{S}} V_{x,y}(\mathbf{q}_n(s))dP_{x,y}$, and therefore

$$\int_{\mathcal{S}} V_{x,y}(\mathbf{p}(s))dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(\mathbf{q}(s))dP_{x,y}.$$

For the other direction, we start with $\int_{\mathcal{S}} V_{x,y}(\mathbf{p}(s))dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(\mathbf{q}(s))dP_{x,y}$. The reversion of the whole argument allows us to conclude that, for any two sequences of rational convergent vectors $\mathbf{p}_n$ converging to $\mathbf{p}$ and $\mathbf{q}_n$ converging to $\mathbf{q}$, it must be that,

$$\forall a, b \in \left(V^{-1}_{x,y}\left(\int_{\mathcal{S}} V_{x,y}(\mathbf{q}(s))dP_{x,y} \right), V^{-1}_{x,y}\left(\int_{\mathcal{S}} V_{x,y}(\mathbf{p}(s))dP_{x,y} \right) \right) \subset [0,1],$$

we have

$$\mathbf{p}_n \succ_{x,y} \mathbf{a} \succ_{x,y} \mathbf{b} \succ_{x,y} \mathbf{q}_n$$

for all large n. We hence have $lim_{n\to\infty} m_{x,y}(\mathbf{p}_n) > lim_{n\to\infty} m_{x,y}(\mathbf{q}_n)$, and therefore $\mathbf{p}\bar{\succ}_{x,y}\mathbf{q}$ by the definition of $\bar{\succcurlyeq}_{x,y}$.

Step 1.5: Now we map the representation of $\bar{\succcurlyeq}_{x,y}$ on $[0,1]^{\mathcal{S}}$ back to the preference $\succcurlyeq$

[11] We can also use the fact that $\{s'\}^c$ is essential. Here we have $m_{x,y}(\mathbf{a}_{s'}\mathbf{p}'^{\mathbb{Q}}) < m_{x,y}(\mathbf{b}_{s'}\mathbf{q}'^{\mathbb{Q}})$ and we can therefore find rational $\mathbf{a}^{\mathbb{Q}} > \mathbf{a}, \mathbf{b}^{\mathbb{Q}} < \mathbf{b}$ that satisfy both $m_{x,y}(\mathbf{a}^{\mathbb{Q}}_s\mathbf{p}^{\mathbb{Q}}) > m_{x,y}(\mathbf{b}^{\mathbb{Q}}_s\mathbf{q}^{\mathbb{Q}})$ and $m_{x,y}(\mathbf{a}^{\mathbb{Q}}_{s'}\mathbf{p}'^{\mathbb{Q}}) < m_{x,y}(\mathbf{b}^{\mathbb{Q}}_{s'}\mathbf{q}'^{\mathbb{Q}})$.

over acts. Specifically, for any acts $f, g \in \mathcal{F}$ with $x \succcurlyeq f(s) \succcurlyeq y, x \succcurlyeq g(s) \succcurlyeq y$ for all $s \in \mathcal{S}$, we claim that

$$f \succcurlyeq g \Leftrightarrow \int_{\mathcal{S}} V_{x,y}(m_{x,y}(f(s)))dP_{x,y} \geq \int_{\mathcal{S}} V_{x,y}(m_{x,y}(g(s)))dP_{x,y}. \tag{4.3}$$

That is, acts f and g's *profile of matching frequency* — $\mathbf{p}_f = (m_{x,y}(f(s)))_{s\in\mathcal{S}}$ and $\mathbf{p}_g = (m_{x,y}(g(s)))_{s\in\mathcal{S}}$, respectively — represent the acts in the sense that $f \succcurlyeq g \Leftrightarrow \mathbf{p}_f \bar{\succcurlyeq}_{x,y} \mathbf{p}_g$. Notice that f, g are not necessarily binary acts, and here and after we are using the notion of matching frequency of general act as in (4.1).

To establish the equivalence (4.2), we start with assuming $f \succ g$. By Lemma 7, we have $m_{x,y}(f) > m_{x,y}(g)$. Let us fix rational numbers $a, b \in (m_{x,y}(g), m_{x,y}(f))$ with $a > b$, and we aim to show $\int_{\mathcal{S}} V_{x,y}(m_{x,y}(f(s)))dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(a)dP_{x,y}$, whist the proof for $\int_{\mathcal{S}} V_{x,y}(m_{x,y}(g(s)))dP_{x,y} < \int_{\mathcal{S}} V_{x,y}(b)dP_{x,y}$ is analogous. Take classical lotteries $\mathbf{x}, \mathbf{y} \in \mathcal{X}_{x,y}$ with $\mathbf{p_x} = \mathbf{a}$ and $\mathbf{p_y} = \mathbf{b}$, and we have $f \succ \mathbf{x} \succ \mathbf{y} \succ g$. By the monotonicity axiom, acts f can be approximated by sequence $\{\mathbf{p}_{f,n}\}_{n\geq 1}$ of rational vectors in $[0,1]^{\mathcal{S}}$, with limit $\mathbf{p}_f = (m_{x,y}(f(s)))_{s\in\mathcal{S}}$. Because the additive model (4.2) is continuous, $\int_{\mathcal{S}} V_{x,y}(\mathbf{p}_{f,n}(s))dP_{x,y}$ converges to $\int_{\mathcal{S}} V_{x,y}(m_{x,y}(f(s)))dP_{x,y}$. Fix a rational number $c \in (a, m_{x,y}(f))$ and classical lottery $\mathbf{z} \in \mathcal{X}_{x,y}$ with $\mathbf{p_z} = \mathbf{c}$, and we have $f \succ \mathbf{z} \succ \mathbf{x}$. If $\int_{\mathcal{S}} V_{x,y}(m_{x,y}(f(s)))dP_{x,y} \leq \int_{\mathcal{S}} V_{x,y}(a)dP_{x,y}$, any rational sequence $\mathbf{p}_{f,n}$ that approaches $\mathbf{p}_f = (m_{x,y}(f(s)))_{s\in\mathcal{S}}$ from above (that is, $\mathbf{p}_{f,n}(s) \geq m_{x,y}(f(s))$ for all state s and all n) shall have $\int_{\mathcal{S}} V_{x,y}(\mathbf{p}_{f,n}(s))dP_{x,y} < \int_{\mathcal{S}} V_{x,y}(c)dP_{x,y}$ for all large n, which in addition to monotonicity imply $\mathbf{z} \succ f_n \succcurlyeq f$ for large n, where $f_n \in \mathcal{F}_{x,y}$ is any act that satisfies $\mathbf{p}_{f_n} = \mathbf{p}_{f,n}$. This contradicts $f \succ \mathbf{z}$.

For the other direction, suppose that $\int_{\mathcal{S}} V_{x,y}(m_{x,y}(f(s)))dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(m_{x,y}(g(s)))dP_{x,y}$. For any sequence $\mathbf{p}_{f,n}$ of rational vectors that approaches $\mathbf{p}_f = (m_{x,y}(f(s)))_{s\in\mathcal{S}}$ from below (that is, $\mathbf{p}_{f,n}(s) \leq m_{x,y}(f(s))$ for all state s and all n) and any sequence of rational vectors $\mathbf{p}_{g,n}$ that approaches $\mathbf{p}_g = (m_{x,y}(g(s)))_{s\in\mathcal{S}}$ from above, by continuity of the additive model (4.2) we shall have $\int_{\mathcal{S}} V_{x,y}(\mathbf{p}_{f,n}(s))dP_{x,y} > \int_{\mathcal{S}} V_{x,y}(\mathbf{p}_{g,n}(s))dP_{x,y}$ for all large n and hence $\mathbf{p}_{f,n} \succ_{x,y} \mathbf{p}_{g,n}$ for all large n. For all n, monotonicity implies $f \succcurlyeq f_n$ and $g_n \succcurlyeq g$ for any acts $f_n, g_n \in \mathcal{F}_{x,y}$ that satisfy $\mathbf{p}_{f_n} = \mathbf{p}_{f,n}$ and $\mathbf{p}_{g_n} = \mathbf{p}_{g,n}$. Hence, we have $f \succcurlyeq f_n \succ g_n \succcurlyeq g$ for all large n, and therefore $f \succ g$.

Step 2.1: From now on we fix a pair of prizes $x \succ y$, and consider any other pair of prizes $x' \succ y'$ with $x' \succcurlyeq x \succ y \succcurlyeq y'$. Previous steps of the proof also hold for the prize pair $x' \succ y'$. Specifically, as in model (4.2, 4.3), we get a continuous, strictly increasing function $V_{x',y'} : [0,1] \to \mathbb{R}$ and a probability measure $P_{x',y'}$ on $\mathcal{S}$ such that

$$f \succcurlyeq g \Leftrightarrow \int_{\mathcal{S}} V_{x',y'}(m_{x',y'}(f(s)))dP_{x',y'} \geq \int_{\mathcal{S}} V_{x',y'}(m_{x',y'}(g(s)))dP_{x',y'}. \tag{4.4}$$

for all acts $f, g \in \mathcal{F}$ with $x' \succcurlyeq f(s) \succcurlyeq y', x' \succcurlyeq g(s) \succcurlyeq y'$ for all $s \in \mathcal{S}$. This model also applies to $\mathcal{F}_{x,y}$ because for all $f \in \mathcal{F}_{x,y}$ and $s \in \mathcal{S}$ we have $x' \succcurlyeq f(s) \succcurlyeq y'$. Moreover, the two models $V_{x,y}$ and $V_{x',y'}$ are related in the following sense. For all $\mathbf{p}, \mathbf{q} \in [0,1]^{\mathcal{S}}$, considered as vectors in the domain of $\bar{\succcurlyeq}_{x,y}$, let $\mathbf{p}', \mathbf{q}' \in [m_{x',y'}(y), m_{x',y'}(x)]^{\mathcal{S}}$ be the affine-transformed vectors (in the domain of $\bar{\succcurlyeq}_{x',y'}$) that, for all $s \in \mathcal{S}$,

$$\begin{aligned} \mathbf{p}'(s) &= m_{x',y'}(y) + \mathbf{p}(s)(m_{x',y'}(x) - m_{x',y'}(y)), \\ \mathbf{q}'(s) &= m_{x',y'}(y) + \mathbf{q}(s)(m_{x',y'}(x) - m_{x',y'}(y)), \end{aligned}$$

and, we claim that

$$\mathbf{p}\bar{\succcurlyeq}_{x,y}\mathbf{q} \Leftrightarrow \int_{\mathcal{S}} V_{x',y'}(\mathbf{p}'(s))dP_{x',y'} \geq \int_{\mathcal{S}} V_{x',y'}(\mathbf{q}'(s))dP_{x',y'} \Leftrightarrow \mathbf{p}'\bar{\succcurlyeq}_{x',y'}\mathbf{q}'. \tag{4.5}$$

The reason is as follows. Relation (4.5) holds for all rational vectors $\mathbf{p}, \mathbf{q} \in [0,1]^{\mathcal{S}}$ because, first, rational vectors $\mathbf{p}$ and $\mathbf{q}$ stand for two acts $f, g \in \mathcal{F}_{x,y}$, second, $\mathbf{p}', \mathbf{q}' \in [m_{x',y'}(y), m_{x',y'}(x)]^{\mathcal{S}}$ essentially stand for the same two acts but now expressed using prize pair $\{x', y'\}$ (i.e., $\mathbf{p}'(s) = m_{x',y'}(f(s))$ and $\mathbf{q}'(s) = m_{x',y'}(g(s))$ for all s) and, third, the relation (4.4) holds. More generally, since $\bar{\succcurlyeq}_{x,y}$ is a continuous extension of $\succcurlyeq_{x,y}$ from $\mathbb{Q}^{\mathcal{S}}_{[0,1]}$ to $[0,1]^{\mathcal{S}}$ as in (4.2) and also because the model (4.4) is continuous, (4.5) has to hold for all $\mathbf{p}, \mathbf{q} \in [0,1]^{\mathcal{S}}$ if it holds for the rational ones. By the uniqueness property of additive representation, (4.2) and (4.5) together imply that

$$\begin{aligned} P_{x,y} &= P_{x',y'}, \\ V_{x,y}(q) &= aV_{x',y'}(m_{x',y'}(y) + q(m_{x',y'}(x) - m_{x',y'}(y))) + b, \end{aligned}$$

for some $a > 0$, some $b \in \mathbb{R}$, and all $q \in [0,1]$. Since the average verifiable utility model implies that for all $q \in [0,1]$,

$$m_{x',y'}(y) + q\left(m_{x',y'}(x) - m_{x',y'}(y)\right) = \frac{qu(x) + (1-q)u(y) - u(y')}{u(x') - u(y')},$$

the same equation can also be written as,

$$V_{x,y}(q) = aV_{x',y'}\left(\frac{qu(x) + (1-q)u(y) - u(y')}{u(x') - u(y')}\right) + b. \tag{4.6}$$

From now on, we denote the common probability measure by $P = P_{x,y}$, and, after a possible affine transformation of $V_{x',y'}$, we take $a = 1$ and $b = 0$. That is, we have

$$\begin{aligned} P &= P_{x',y'}, \\ V_{x,y}(q) &= V_{x',y'}\left(\frac{qu(x) + (1-q)u(y) - u(y')}{u(x') - u(y')},\right) \end{aligned} \tag{4.7}$$

for all $q \in [0,1]$ and all $x' \succ y'$ with $x' \succcurlyeq x \succ y \succcurlyeq y'$.

Step 2.2: With relation (4.7), we can patch the models on different prize pairs together to form a unified model. We start by normalizing the average verifiable utility $U : \mathcal{X} \to \mathbb{R}$

with $u(x) = 1$ and $u(y) = 0$, where $x \succ y$ are the two (albeit arbitrary) prizes we fixed before. Since for any classical lottery $\mathbf{z}$ with $x \succcurlyeq \mathbf{z} \succcurlyeq y$ we have $m_{x,y}(\mathbf{z}) = U(z)$, model (4.3) can be rewritten as

$$f \succcurlyeq g \Leftrightarrow \int_{\mathcal{S}} V_{x,y}(U(f(s)))dP \geq \int_{\mathcal{S}} V_{x,y}(U(g(s)))dP,$$

for all acts $f, g \in \mathcal{F}$ with $x \succcurlyeq f(s) \succcurlyeq y, x \succcurlyeq g(s) \succcurlyeq y$ for all $s \in \mathcal{S}$. This is simply the smooth model applied to the set of acts $\{f \in \mathcal{F} : x \succcurlyeq f(s) \succcurlyeq y, \forall s\}$. Hence we can define function $\psi : conv\,(U(\mathcal{X})) \to \mathbb{R}$ first on $[0,1]$ by $\psi(e) = V_{x,y}(e)$ for all $e \in [0,1]$. For any other pair of prizes $x' \succ y'$ with $x' \succcurlyeq x \succ y \succcurlyeq y'$, we extend ψ to $[u(y'), u(x')]$ by setting

$$\psi(e) = V_{x',y'}\left(\frac{e - u(y')}{u(x') - u(y')}\right), \forall e \in [u(y'), u(x')]. \tag{4.8}$$

For ψ to be well-defined, we need to show that, if for some $e^* \in [u(y'), u(x')]$ there is another pair of prizes $x'' \succ y''$ with $x'' \succcurlyeq x \succ y \succcurlyeq y''$ and $e^* \in [u(y''), u(x'')]$, then we have

$$V_{x',y'}\left(\frac{e^* - u(y')}{u(x') - u(y')}\right) = V_{x'',y''}\left(\frac{e^* - u(y'')}{u(x'') - u(y'')}\right). \tag{4.9}$$

To see why this is indeed the case, let x^* denote the more preferred prize between x' and x'' and let y^* denote the less preferred prize between y' and y''. Following the same arguments leading to equation (4.6), we are able to obtain the following equations that, for some $a_1, a_2 > 0$ and some $b_1, b_2 \in \mathbb{R}$,

$$\begin{aligned}
V_{x',y'}\left(\frac{e_1 - u(y')}{u(x') - u(y')}\right) &= a_1 V_{x^*,y^*}\left(\frac{e_1 - u(y^*)}{u(x^*) - u(y^*)}\right) + b_1, \\
V_{x'',y''}\left(\frac{e_2 - u(y'')}{u(x'') - u(y'')}\right) &= a_2 V_{x^*,y^*}\left(\frac{e_2 - u(y^*)}{u(x^*) - u(y^*)}\right) + b_2,
\end{aligned}$$

for all $e_1 \in [u(y'), u(x')]$ and all $e_2 \in [u(y''), u(x'')]$. But by equation (4.7) we have that, for all $e \in [0,1]$,

$$V_{x,y}(e) = V_{x',y'}\left(\frac{e - u(y')}{u(x') - u(y')}\right) = V_{x'',y''}\left(\frac{e - u(y'')}{u(x'') - u(y'')}\right) = V_{x^*,y^*}\left(\frac{e - u(y^*)}{u(x') - u(y')}\right),$$

which implies $a_1 = a_2 = 1$ and $b_1 = b_2 = 0$. This validates equation (4.9). After taking into account all prize pairs $x' \succ y'$ with $x' \succcurlyeq x \succ y \succcurlyeq y'$, we obtain a well-defined function $\psi : conv\,(U(\mathcal{X})) \to \mathbb{R}$.

Step 2.3: Finally we show that the smooth model with probability measure P and function ψ does represent the preference. For any two acts $f, g \in \mathcal{F}$, fix two prizes $x' \succ y'$ with $x' \succcurlyeq x \succ y \succcurlyeq y'$ and $x' \succcurlyeq f(s) \succcurlyeq y'$ and $x' \succcurlyeq g(s) \succcurlyeq y'$ for all s. For example, we can take x' to be the better prize between x and the best prize in the support of f and g, and y' to be the lesser prize between y and the worst prize in the support of f and g. By definition

(4.8), $\psi(U(f(s))) = V_{x',y'}\left(\frac{U(f(s))-u(y')}{u(x')-u(y')}\right)$ for all s, and, because $\frac{U(f(s))-u(y')}{u(x')-u(y')} = m_{x',y'}(f(s))$ for all s, we have

$$\int_{\mathcal{S}} \psi(U(f(s)))dP = \int_{\mathcal{S}} V_{x',y'}(m_{x',y'}(f(s)))dP.$$

Similarly, we have $\int_{\mathcal{S}} \psi(U(g(s)))dP = \int_{\mathcal{S}} V_{x',y'}(m_{x',y'}(g(s)))dP$. By (4.4) we can conclude that

$$f \succcurlyeq g \Leftrightarrow \int_{\mathcal{S}} \psi(U(f(s)))dP \geq \int_{\mathcal{S}} \psi(U(g(s)))dP.$$

The proof is hence complete. □

4.4 Proof of Proposition 6

Proof. We start with unverifiable uncertainty aversion and, because its necessity is straightforward, we shall focus on the sufficiency part. The case of unverifiable uncertainty seeking is symmetric and will be omitted.

As in the Step 1.1 of Proof of Theorem 4, we obtain for any two prizes $x \succ y$ a weak-order preference $\succcurlyeq_{x,y}$ on $\mathbb{Q}_{[0,1]}^{\mathcal{S}}$, which is derived from preference $\succcurlyeq$ over the set $\mathcal{F}_{x,y}$ of binary acts supported on $\{x, y\}$. Using the construction in the proof of Theorem 4, the model (4.2) represents $\succcurlyeq_{x,y}$ on $\mathbb{Q}_{[0,1]}^{\mathcal{S}}$ and, actually, its extended version $\bar{\succcurlyeq}_{x,y}$ on the full domain $[0,1]^{\mathcal{S}}$.

Given that $\succcurlyeq_{x,y}$ on $\mathbb{Q}_{[0,1]}^{\mathcal{S}}$ exhibits unverifiable uncertainty aversion, we claim that $\bar{\succcurlyeq}_{x,y}$ is a convex preference, that is, for all $\mathbf{q} \in [0,1]^{\mathcal{S}}$ the set $\{\mathbf{p} \in [0,1]^{\mathcal{S}} : \mathbf{p} \bar{\succcurlyeq} \mathbf{q}\}$ is convex. If this is not true, then there exist $\mathbf{p}, \mathbf{q} \in [0,1]^{\mathcal{S}}$ with $\mathbf{p} \bar{\succcurlyeq}_{x,y} \mathbf{q}$ and $\mathbf{q} \bar{\succ}_{x,y} \alpha\mathbf{p} + (1-\alpha)\mathbf{q}$ for some $\alpha \in (0,1)$. Find any sequences $\mathbf{p}_n, \mathbf{q}_n \in \mathbb{Q}_{[0,1]}^{\mathcal{S}}$ and $\alpha_n \in \mathbb{Q}_{[0,1]}$ such that $\mathbf{p}_n$ converges to $\mathbf{p}$, $\mathbf{q}_n$ converges to $\mathbf{q}$, and α_n to α. By continuity, we have $\mathbf{p}_n \bar{\succ}_{x,y} \alpha_n\mathbf{p} + (1-\alpha_n)\mathbf{q}_n$ and $\mathbf{q}_n \bar{\succ}_{x,y} \alpha_n\mathbf{p} + (1-\alpha_n)\mathbf{q}_n$ for all large n, which contradicts unverifiable uncertainty aversion for all these large n (no matter $\mathbf{p}_n \bar{\succcurlyeq}_{x,y} \mathbf{q}_n$ or $\mathbf{q}_n \bar{\succcurlyeq}_{x,y} \mathbf{p}_n$).

By Debreu and Koopmans (1982), the function $V_{x,y}$ of the model (4.2) is concave on $[0,1]$, and, therefore, the function ψ is concave on $[U(y), U(x)]$. Since the choice of x and y is arbitrary, we can conclude that ψ is concave on its whole domain. □