9-16-2024

Generalized NEO-EU Preferences and the Falsifiability of Ambiguity Theories

Manuel Nunez

Mark Schneider

Generalized NEO-EU Preferences and the Falsifiability of Ambiguity Theories

Comments

ESI Working Paper 24-15

Generalized NEO-EU Preferences and the Falsifiability of Ambiguity Theories

Manuel Nunez and Mark Schneider

September 16, 2024

Abstract

Axiomatic non-expected utility models are generally more difficult to falsify than expected utility theory as they are less restrictive (by weakening the independence axiom). Recent work computes the Vapnik-Chervonenkis (VC) dimension of a theory to determine the extent to which the theory is falsifiable. Popular ambiguity theories have VC dimensions that increase exponentially in the number of states or that are infinite, whereas the VC dimension of expected utility theory increases linearly in the number of states. In this paper we axiomatically characterize the class of generalized non-extreme outcome expected utility (NEO-EU) preferences in the Anscombe-Aumann framework and show that their VC dimension increases linearly in the number of states. Our paper shows that this popular class of ambiguity preferences which has been broadly applied provides a counter-example to the conjecture that axiomatic models of ambiguity attitudes are substantially more difficult to falsify than expected utility theory.

JEL classification: D81.

Keywords: Generalized NEO-EU; Choice under Ambiguity; Decision Theory

*Affiliation of authors: Manuel Nunez, University of Connecticut Business, University of Connecticut, 2100 Hillside Rd, Storrs, CT 06269. Electronic address: `manuel.nunez@uconn.edu`. Phone: (860) 486-1333; Mark Schneider, Culverhouse College of Business, University of Alabama, 361 Stadium Drive, Tuscaloosa, AL 35487. Electronic address: `maschneider4@ua.edu`. Phone: (205) 348-1554.

1 Introduction

Over a century ago, Knight (1921) and Keynes (1921) recognized that situations of ambiguity, in which probabilities of events are unknown, play a fundamental role in decisions and markets. Subjective expected utility (SEU) theory (Savage, 1954; Anscombe and Aumann, 1963), the workhorse model in theoretical and empirical economic applications, can only describe neutral attitudes toward ambiguity, a behavior violated empirically (Ellsberg, 1961). As noted by Baillon et al. (2018b), Ellsberg's paradox "showed that fundamentally new models are needed to handle ambiguity." However, it has been shown in a precise sense, that while SEU is in principle falsifiable, popular axiomatic ambiguity models are substantially more difficult or are even impossible to falsify (Basu and Echenique, 2020). It is an open question whether all popular models of ambiguity preferences are substantially more difficult to falsify than SEU or whether a popular class of axiomatic ambiguity models is essentially no more difficult to falsify than SEU.

In this paper we study the foundations of a class of ambiguity preferences that have seen increasing and broad applications in economics: the class of generalized non-extreme outcome expected utility (NEO-EU) preferences (Chateauneuf et al., 2007; Eichberger et al., 2012). In particular, we study the axiomatic foundations and the complexity of generalized NEO-EU preferences, and establish three main results.

First, we introduce a new weakening of the SEU independence axiom, *co-extreme independence*, that imposes independence only between pairs of acts whose best payoffs occur in the same state and whose worst payoffs occur in the same state. The axiom implies that independence holds between pairs of acts for which hedging cannot change the ranking of which act is more secure or which act has more potential (in the sense of Frick et al. (2022)). However, the axiom permits violations of independence for acts where hedging can change which act is more secure or which act has more potential. Our first main result is that together with the standard invariant bi-separable axioms (Ghirardato et al., 2004), co-extreme independence characterizes the class of generalized NEO-EU preferences.

Second, we consider the complexity and falsifiability of generalized NEO-EU preferences. Basu and Echenique (2020), measure the degree of falsifiability of a decision theory by computing its Vapnik-Chervonenkis (VC) dimension, which is the largest sample size for which the theory can always rationalize the data. As observed by

Chambers et al. (2023), the VC dimension "is a measure of the complexity of a theory used in machine learning." Basu and Echenique (2020) show that the VC dimension increases linearly in the number of states for SEU, but it increases exponentially in the number of states for Choquet expected utility theory (Schmeidler, 1989), which weakens SEU independence to co-monotonic independence, and it is infinite when there are more than two states for the multiple priors model (Gilboa and Schmeidler, 1989; Ghirardato et al., 2004), which weakens SEU independence to certainty independence. Popular axiomatic ambiguity models that generalize the multiple priors model such as invariant bi-separable preferences (Ghirardato et al., 2004), variational preferences (Maccheroni et al., 2006), mean-dispersion preferences (Grant and Polak, 2013), and $\alpha-$maxmin preferences (Hartmann, 2023) likewise have infinite VC dimension when there are more than two states. Basu and Echenique (2020) conclude that SEU is in principle falsifiable, while the popular ambiguity models appear difficult or even impossible to falsify. Although co-extreme independence is much less restrictive than SEU independence as it imposes independence for a considerably smaller class of acts, our second main result is that the VC dimension of generalized NEO-EU preferences (and of NEO-EU preferences) increases linearly in the number of states, similar to SEU. That is, NEO-EU preferences are essentially no more difficult to falsify in principle than SEU and the two models have similar complexity.

Our third main result is that we introduce an algorithm for identifying and computing the NEO-EU parameters from an arbitrary finite dataset of binary choices. We provide conditions under which this algorithm runs in polynomial time, and we observe that it has essentially the same computational complexity as an algorithm designed to identify and compute the parameters of SEU. This result complements the findings in Chambers et al. (2021) and Chambers et al. (2023) who study, respectively, a general problem of recovering preferences and utility functions from finite datasets, but do not introduce polynomial time algorithms for computing the preferences.

The class of generalized NEO-EU preferences includes several popular and familiar ambiguity models as special cases including the $\epsilon-$contamination model (Dow and da Costa Werlang, 1992), the NEO-EU model (Chateauneuf et al., 2007), Hurwicz preferences (Hurwicz, 1950; Grant and Polak, 2013), and maxmin preferences (Wald, 1950) as well as SEU.[1] The special case of NEO-EU preferences allows for a broad

[1]Eichberger et al. (2012) formulate generalized NEO-EU preferences in terms of non-additive probabilities or capacities and study them in a sequential setting of purely subjective uncertainty.

range of ambiguity attitudes, ranging from extreme ambiguity aversion to extreme optimism toward ambiguity, and it achieves a separation of preferences into a parameter representing the agent's ambiguity attitude and a parameter representing the agent's perceptions of ambiguity.[2].

Our paper contributes to the understanding of the foundations and properties of the NEO-EU model (Chateauneuf et al., 2007; Eichberger et al., 2012), which provides a compromise between SEU and the more general ambiguity models.[3] The NEO-EU model is perhaps the simplest generalization of SEU that explains two of the most robust features of observed ambiguity attitudes: (i) an aversion toward ambiguity for symmetric and negatively skewed ambiguous payoffs (Ellsberg, 1961; Baillon and Bleichrodt, 2015; Dimmock et al., 2015; Kocher et al., 2018); and (ii) a preference for ambiguity toward positively skewed ambiguous payoffs (Baillon and Bleichrodt, 2015; Dimmock et al., 2015; Kocher et al., 2018).

While many of the major ambiguity models have an established axiomatic characterization in the standard Anscombe and Aumann (1963) framework, NEO-EU and generalized NEO-EU lack an established characterization in the Anscombe-Aumann framework.[4] In addition to characterizing generalized NEO-EU preferences in the Anscombe-Aumann framework, we also introduce a weakening of the "preference for hedging" axiom that helps characterize the class of *ambiguity-averse NEO-EU preferences* (the class of NEO-EU preferences consistent with ambiguity aversion in Ellsberg's paradox) in the Anscombe-Aumann framework. This axiom restricts the simple diversification axiom in Siniscalchi (2009) to a subset of complementary acts (Siniscalchi, 2009; Chambers et al., 2014), and it does not conflict with ambiguity

They show that convex generalized NEO-EU capacities are "the only capacity for which the core of the full Bayesian updates of a capacity equals the set of Bayesian updates of the probability distributions in the core of the original capacity."

[2]The ambiguity parameter represents ambiguity in a precise sense as it controls the size of the agent's set of prior distributions under the multiple priors representation of the NEO-EU model

[3] Wakker (2010) notes that NEO-EU is among the more promising non-expected utility models and that the interpretation of its parameters is clearer and more convincing than with other models. He also comments that NEO-EU may reflect an optimal trade-off between parsimony and fit.

[4]Major ambiguity models that have an established axiomatic characterization in the Anscombe-Aumann framework include Choquet expected utility theory (Schmeidler, 1989), the maxmin multiple priors model (Gilboa and Schmeidler, 1989), invariant bi-separable preferences (Ghirardato et al., 2004), variational preferences (Maccheroni et al., 2006), vector expected utility (Siniscalchi, 2009), the smooth model of ambiguity aversion (Neilson, 2010; Denti and Pomatto, 2022), multiplier preferences (Strzalecki, 2011), mean-dispersion preferences (Grant and Polak, 2013), and α-maxmin preferences (Hartmann, 2023).

seeking behavior toward low-likelihood gains which is ruled out by the standard preference for hedging axiom.

The importance of better understanding the axiomatic foundations of NEO-EU is further underscored by the growing number of applications of NEO-EU to diverse economic problems. For instance, the NEO-EU model has been applied to explain buying-selling price gaps in markets (Dow and da Costa Werlang, 1992), to reconcile the simultaneous purchasing of lottery tickets and insurance policies (Chateauneuf et al., 2007), to generalize the classical Consumption CAPM (Zimper, 2012), to generalize the life-cycle model of consumption and saving (Groneck et al., 2016), to study the design of optimal contracts under asymmetric information (Giraud and Thomas, 2017), to explain demand elasticities from state-run lotteries (Lockwood et al., 2021), to explain time variation in the stock market risk-return tradeoff (Ghazi et al., 2023a), and to explain momentum in stock returns (Ghazi et al., 2023b).

The remainder of the paper is organized as follows: Section 2 studies the axiomatic foundations of NEO-EU. Section 3 studies the falsifiability of NEO-EU and presents an algorithm for the identification of the NEO-EU parameters. Section 4 concludes.

2 NEO-EU Preferences

In this section we use standard axioms and propose new axioms to help characterize the class of generalized NEO-EU preferences and the class of ambiguity-averse NEO-EU preferences.

2.1 Preliminaries

We work in the Anscombe and Aumann (1963) framework in which there is both objective and subjective uncertainty. Let X denote a finite set of outcomes with at least two elements.[5] An *objective lottery*, $p : X \to [0, 1]$, is a probability distribution over outcomes. Let $\Delta(X)$ denote the set of objective lotteries and assume that it is a mixture space. A *von Neumann-Morgenstern (vNM) expected utility function* is an

[5] Assuming that X is finite is without loss of generality because, as standard in the literature, once a representation has been established for a finite set of outcomes it is not difficult to extend it to an infinite countable set, which is sufficient for most practical applications.

application $U : \Delta(X) \to \mathbb{R}$ defined as

$$U(p) := \sum_{x \in X} p(x)u(x),$$

where $u : X \to \mathbb{R}$ is an arbitrary function defined on the set of outcomes (we refer to this function simply as the *utility* function).

Let S be a totally ordered finite set with at least two elements (this is to avoid trivial situations) indexing all states of nature. We define a *subjective lottery* or *act* f as any mapping $f : S \to \Delta(X)$. The set of probability distributions on S is denoted by $\Delta(S)$. We will use $f(s)$, $s \in S$, to denote the objective lottery assigned by f to state s. We denote the set of acts by $\mathcal{F}$ and assume that it is a mixture space. We define a *constant* act f as an act that yields the same objective lottery in every state of nature, i.e., $f(s) = p$ for all $s \in S$, where $p \in \Delta(X)$ is an objective lottery. In this case, abusing notation, we also let p denote the corresponding constant act and write $p \in \mathcal{F}$ and $\Delta(X) \subset \mathcal{F}$.

We assume that there is a binary relation denoted by "$\succ$" $\subset \mathcal{F} \times \mathcal{F}$ over $\mathcal{F}$. The relation $\succ$ is called a *preference relation* if it is asymmetric and negatively transitive, and in that case, we say that f is *preferred to* g if $f \succ g$. Moreover, we say that f is *weakly preferred* to g, denoted as $f \succsim g$, if $g \not\succ f$; and that f is *indifferent* to g, denoted as $f \sim g$, if $f \not\succ g$ and $g \not\succ f$. Observe that if $\succ$ is a preference relation, then for all f and g exactly one of $f \succ g$, $g \succ f$, or $f \sim g$ holds; and $\succsim$ is a complete and transitive relation (Kreps, 1988). Using constant acts, we can then extend a preference relation to $\Delta(X)$ by writing $p \succ q$, for $p, q \in \Delta(X)$, whenever the constant act yielding lottery p is preferred to the constant act yielding lottery q. Similarly, we can extend a preference relation to X by writing $x \succ y$, for $x, y \in X$, whenever the constant act yielding lottery p that assigns probability one to x is preferred to the constant act yielding lottery q that assigns probability one to y. Since X is finite, we say that $\succ$ is nontrivial if there exist $\overline{x}, \underline{x} \in X$ such that $\overline{x} \succsim x \succsim \underline{x}$ for all $x \in X$, and $\overline{x} \succ \underline{x}$.

2.2 Invariant Bi-Separable and NEO-EU Preferences

We state standard axioms that summarize basic assumptions about the $\succ$ relation:

Axiom 1 *(Preference) $\succ$ on $\mathcal{F}$ is a nontrivial preference relation.*

Axiom 2 *(Continuity) For every $f, g, h \in \mathcal{F}$, the sets $\{\gamma \in [0,1] : \gamma f + (1-\gamma)g \succsim h\}$ and $\{\gamma \in [0,1] : h \succsim \gamma f + (1-\gamma)g\}$ are closed.*

Axiom 3 *(Certainty independence) For every $f, g \in \mathcal{F}$, $p \in \Delta(X)$ and $\gamma \in (0,1]$, we have $f \succsim g$ if and only if $\gamma f + (1-\gamma)p \succsim \gamma g + (1-\gamma)p$.*

Axiom 3 is introduced in Gilboa and Schmeidler (1989) in their characterization of the maxmin multiple priors model, and it is the key axiom used to characterize invariant bi-separable preferences in Ghirardato et al. (2004). Axioms 1 to 3 imply the classical von Neumann-Morgenstern representation theorem when restricted to the set $\Delta(X)$ of objective lotteries (Kreps, 1988):

Theorem 1 *Axioms 1 to 3 are necessary and sufficient for there exists a nonconstant function $u : X \to \mathbb{R}$ such that*

$$p \succ q \text{ if and only if } U(p) > U(q). \tag{1}$$

for all $p, q \in \Delta(X)$. Furthermore, $u' : X \to \mathbb{R}$ is a function also representing $\succ$ in the sense of (1) if and only if there exist real numbers $c > 0$ and d such that $u'(x) = cu(x) + d$ for all $x \in X$.

From now on, we assume without loss of generality that the utility function u from Theorem 1 has been normalized so that

$$u(\overline{x}) = 1, u(\underline{x}) = 0. \tag{2}$$

This implies that

$$0 \leq u(x) \leq 1 \text{ and } 0 \leq U(p) \leq 1, \tag{3}$$

for all $x \in X$ and $p \in \Delta(X)$.

Axiom 4 *(Monotonicity) For every $f, g \in \mathcal{F}$, $f(s) \succsim g(s)$ for all $s \in S$ implies $f \succsim g$.*

We next summarize important classes of preferences that satisfy Axioms 1 through 4:

1. **Subjective Expected Utility Preferences (SEU):** Given a distribution $\pi \in \Delta(S)$, the *expected utility* functional on $\mathcal{F}$ with respect to π is defined as

$$\mu(f;\pi) := \sum_{s\in S} \pi_s U(f(s)). \tag{4}$$

   When the distribution π is clear from the context, we will simply write $\mu(f)$ for the corresponding functional.

2. **Invariant Biseparable Preferences (IB):** Ghirardato et al. (2004) show that Axioms 1 to 4 imply the existence of a nonempty, weak* compact convex set $\mathcal{P} \subset \Delta(S)$, a nonconstant function $u : X \to \mathbb{R}$, and a function $a : \mathcal{F} \to [0,1]$ such that $\succ$ is represented by:

$$I(f) = a(f) \max_{P\in\mathcal{P}} \mu(f;P) + (1 - a(f)) \min_{P\in\mathcal{P}} \mu(f;P). \tag{5}$$

   Further, the set $\mathcal{P}$ is unique and u is unique up to a positive affine transformation.

3. α**-maxmin preferences:** Note that for IB preferences in (5), a depends on the act f and so could differ across acts. An important special case of IB preferences is the class of α-maxmin preferences in which a is a constant denoted by α:

$$I(f) = \alpha \max_{P\in\mathcal{P}} \mu(f;P) + (1 - \alpha) \min_{P\in\mathcal{P}} \mu(f;P). \tag{6}$$

   An appealing feature of α-maxmin preferences is that they provide an approach to separating the decision maker's perceptions of ambiguity (represented by the set of priors $\mathcal{P}$), and the decision maker's attitude toward ambiguity represented by $\alpha \in [0,1]$, which can be interpreted as the agent's degree of ambiguity aversion. These preferences have recently been characterized in the Anscombe-Aumann framework by Hartmann (2023). A notable limitation of α-maxmin preferences, however, as noted in Hartmann (2023), is that neither α nor the set of priors $\mathcal{P}$ is uniquely determined from the axioms. Rather, the axioms imply a continuum of α-maxmin preference representations, each with a potentially different α and set of priors, so that the parameter α cannot be uniquely identified from choice data. However, Hartmann (2023) shows that α is *set identified* in that the bounds on α across different α-maxmin representations can be identified from choice data.

4. **NEO-EU preferences:** While the class of α-maxmin preferences is quite general, this generality comes at a cost. For instance, as we observe in Section 3, α-maxmin preferences are not falsifiable in a sense made precise by Basu and Echenique (2020). Apart from that issue, for empirical applications it is often more useful to work with a specific parameterized set of prior distributions. However, the general α-maxmin model does not provide guidance on which set of priors to use, leaving any chosen set of priors to be somewhat arbitrary. The class of NEO-EU preferences (Chateauneuf et al., 2007) are represented by the functional:

$$V(f) := \rho\mu(f;\pi) + (1-\rho)\left[\lambda \max_{s\in S} U(f(s)) + (1-\lambda)\min_{s\in S} U(f(s))\right], \tag{7}$$

where $\lambda \in [0,1]$ represents the decision maker's degree of optimism toward ambiguity, while $\rho \in [0,1]$ parameterizes the degree of ambiguity perceived by the decision maker. NEO-EU preferences are the special case of (6) in which $\mathcal{P} := \{p \in \Delta(S) : p(E) \geq \rho\pi(E) \text{ for all } E \subset S\}$. The set $\mathcal{P}$ is naturally interpreted as a set of probabilities that lie within an interval of the reference prior π. Further, $1-\rho$ represents the decision maker's perceived ambiguity in the sense that the size of the set of priors expands as ρ decreases. When the agent perceives no ambiguity ($\rho = 1$), $\mathcal{P}$ collapses to a singleton $\{\pi\}$, and the model reduces to SEU. Under maximal ambiguity ($\rho = 0$), $\mathcal{P}$ expands to the entire simplex $\Delta(S)$.

2.3 Characterization of Generalized NEO-EU Preferences

We say that an agent has a *generalized non-extreme outcome expected utility* (abbreviated generalized NEO-EU) preference functional with parameters $(\pi, \lambda, \rho) \in \Delta(S) \times \mathbb{R} \times \mathbb{R}$ if the agent uses the functional $V(f) = V(f;\pi,\lambda,\rho)$ as defined in (7) to compare acts in $\mathcal{F}$, except that the parameters satisfy the following conditions:

$$\begin{aligned} \rho &\geq 0, & (8)\\ \rho\pi_s + (1-\rho)\lambda &\geq 0, \text{ for all } s \in S, & (9)\\ \rho\pi_s + (1-\rho)(1-\lambda) &\geq 0, \text{ for all } s \in S. & (10) \end{aligned}$$

This representation includes five prominent decision models as special cases:

1. NEO-EU preferences as defined by Chateauneuf et al. (2007) if $\lambda, \rho \in [0, 1]$.
2. ϵ-contamination preferences if $\rho = 1 - \epsilon$, where $\epsilon \in [0, 1]$, and $\lambda = 0$.
3. SEU preferences if $\rho = 1$.
4. Hurwicz preferences if $\rho = 0$.
5. Maximin preferences if $\rho = 0$ and $\lambda = 0$.

2.4 Co-Extreme Independence and NEO-EU Preferences

We define more concepts to provide a foundation for a generalized NEO-EU representation.

Definition 1 *Given an act $f \in \mathcal{F}$, the following two subsets of S are called the maximum and minimum sets of states of f, respectively:*

$$\begin{aligned} \overline{S}(f) &:= \{s \in S : f(s) \succsim f(s') \text{ for all } s' \in S\}, \\ \underline{S}(f) &:= \{s \in S : f(s') \succsim f(s) \text{ for all } s' \in S\}. \end{aligned}$$

The elements of $\overline{S}(f)$ are called "maximum" states of f, and the elements of $\underline{S}(f)$ are called "minimum" states of f. We will denote by $R(f)$ the set of states in S that are neither maximum nor minimum of f, that is,

$$R(f) := S \setminus \left(\overline{S}(f) \cup \underline{S}(f)\right).$$

Notice that because $\succ$ is a preference relation and S is finite, the maximum and minimum sets of any act are always nonempty. Also, notice that if s and s' are maximum states of f, then we must have $f(s) \sim f(s')$. Similarly, if s and s' are minimum states of f, then we must have $f(s) \sim f(s')$. Also, there exists a state that is both minimum and maximum for the same act f if and only if $\overline{S}(f) = \underline{S}(f) = S$.

Definition 2 *Two acts $f, g \in \mathcal{F}$ are called "state co-extreme", or simply "co-extreme", which we denote by $f \equiv g$, if they have the same sets of maximum and minimum states. That is, f, g are co-extreme if $\overline{S}(f) = \overline{S}(g)$ and $\underline{S}(f) = \underline{S}(g)$.*

We are ready to state a key new axiom that helps in establishing a generalized NEO-EU preference representation and that bridges the gap between the invariant bi-separable representation in (5) and the representation in (7).

Axiom 5 *(Co-extreme independence) For every f, g, h co-extreme to each other and $\gamma \in (0, 1]$, $f \succ g$ implies $\gamma f + (1-\gamma)h \succ \gamma g + (1-\gamma)h$.*

It is easy to see that if the classical independence axiom holds, then Axiom 5 will also hold, but the converse is not necessarily true. For instance, as we show in Theorem 2 below, preferences represented by NEO-EU functionals with $\rho \neq 1$ satisfy Axiom 5, but being nonlinear across non-co-extreme acts, they cannot satisfy the independence axiom. Also, the preferences in Ellsberg's urn paradoxes are examples which violate the independence axiom but not Axiom 5.

There is a hedging intuition behind Axiom 5. A mixture of two co-extreme acts mixes the worst states in one act with the worst states in another, and mixes the best states in one act with the best states in another. Since the essence of hedging is to mix bad states for one act with good states in another act such that the mixture of the two acts is less exposed to uncertainty, co-extreme acts are not particularly useful in hedging tail ambiguity (ambiguity associated with the extreme outcomes of an act). A decision maker who prefers to hedge tail ambiguity might deviate from the independence axiom and exhibit a preference for hedging for acts that are not co-extreme, but otherwise satisfy the independence axiom. Proposition 5 in the Appendix implies that a mixture of co-extreme acts is also co-extreme with those acts. Hence, rankings of the best and worst outcomes across acts are preserved under mixtures of co-extreme acts. Thus, Axiom 5 imposes independence in cases where hedging cannot change which act has the worst outcome or which act has the best outcome. Formally, in a similar spirit to Kopylov (2009) and Frick et al. (2022), we say an act f is *more secure* than an act g if there exists $s \in S$ such that $f(s') \succsim g(s)$ for all $s' \in S$. Similarly, we say an act f has *more potential* than an act g if there exists $s \in S$ such that $f(s) \succsim g(s')$ for all $s' \in S$. In this precise sense, co-extreme independence restricts the independence axiom to acts where hedging cannot change which act is more secure or which act has more potential. In contrast, co-extreme independence permits violations of independence, for example, if f is more secure or has more potential than g, but the mixture $\gamma g + (1-\gamma)h$ is more secure or has more potential than $\gamma f + (1-\gamma)h$.

The co-extreme independence axiom is reminiscent of the tail independence axiom used by Wakker and Zank (2002) as a key axiom in their characterizations of rank-dependent utility preferences and cumulative prospect theory preferences in choice under risk. Wakker and Zank (2002) note that tail independence "weakens independence somewhat further by considering only maximal or minimal common outcomes." Zank (2007) notes that tail independence is formally equivalent to the axiom called independence of common extremes used by Zank (2007) to characterize a class of social welfare functions. However, to our knowledge, such an axiom is missing from the axioms used to characterize ambiguity models in the Anscombe-Aumann framework.

Chateauneuf et al. (2007) use an axiom of Extreme Event Sensitivity (EES) that formalizes optimism (as an aversion to hedging utility on good events) and pessimism (as a preference for hedging utility on bad events) in their characterization of NEO-EU in a purely subjective setting. An axiom of extreme outcome sensitivity (EOS), which assumes preferences are invariant to changes in common intermediate outcomes is used by Toquebeuf (2016) to characterize generalized NEO-EU in a purely subjective setting as a special case of Choquet expected utility. Eichberger and Kelsey (1999) assume an axiom of extremal independence (EI) in their characterization of E-capacities, which applies independence to acts with a common worst state. Their EI axiom leads to a representation that is a weighted sum of Choquet integrals with respect to an ϵ-contamination representation in which the mixture parameter depends on the particular event. Dominiak and Guerdjikova (2021) assume preferences satisfy extreme event independence (EEI) which seeks to adapt EES to the Anscombe-Aumann framework. The EEI axiom leads to a representation that is a weighted sum of Choquet integrals with respect to a neo-additive capacity in which the parameters λ and ρ depend on the state. A surprising property, in light of Dominiak and Guerdjikova (2021), where the preference parameters can be different for each state, and Ghirardato et al. (2004), where the preference parameters can be different for each act, is that under our representation with the co-extreme independence axiom the resulting preference parameters are constant across states and across acts and along with the agent's subjective prior distribution, they are uniquely determined from the axioms.

The next result determines a set of conditions under which a generalized NEO-EU representation will hold for all acts.

Theorem 2 *Let $|S| > 3$. The relation $\succ$ on $\mathcal{F}$ satisfies Axioms 1 through 5 if and only if there exist $\lambda, \rho \in \mathbb{R}$ and $\pi \in \Delta(S)$ such that*

$$f \succ g \text{ if and only if } V(f) > V(g), \tag{11}$$

for all $f, g \in \mathcal{F}$, where V is the generalized NEO-EU preference functional with parameters (π, λ, ρ) from (7) satisfying conditions (8), (9), and (10). The utility function u in the definition of V is nonconstant and unique up to a positive linear transformation. The scalar ρ is unique. Moreover, if $\rho \neq 1$, then the scalar λ is unique. If $\rho > 0$, then the vector π is unique. Further, if $\rho \neq 1$ and there is at least one state $s \in S$ such that $\pi_s = 0$, then conditions (8)-(10) are equivalent to $0 \leq \rho < 1$ and $0 \leq \lambda \leq 1$.

Proof. See Appendix. ■

Consequently, if there is a state s with $\pi_s = 0$, generalized NEO-EU preferences based on Axioms $1 - 5$ reduce to the standard NEO-EU representation with $\rho, \lambda \in [0, 1]$.

2.5 Additional Axioms and Refinements

In general, the representation from the previous section allows for parameters λ and ρ outside of the interval $[0, 1]$. The next result obtains further restrictions on the parameters by adapting the definition of complementary acts due to Siniscalchi (2009):

Definition 3 *Two acts f and $\hat{f}$ are complementary with indifferent extremes, abbreviated CIE, if $\frac{1}{2}f(s) + \frac{1}{2}\hat{f}(s) \sim \frac{1}{2}f(s') + \frac{1}{2}\hat{f}(s')$ for all $s, s' \in S$, there exist $s \in \overline{S}(f)$ and $s' \in \overline{S}(\hat{f})$ such that $f(s) \sim \hat{f}(s')$, and there exist $s \in \underline{S}(f)$ and $s' \in \underline{S}(\hat{f})$ such that $f(s) \sim \hat{f}(s')$.*

An example of CIE acts is the pair of ambiguous acts in Ellsberg's two-color paradox.

Axiom 6 *(Uncertainty aversion for CIE acts) For each $f, g \in \mathcal{F}$, if f, g is a CIE pair and $f \sim g$ then $\frac{1}{2}f + \frac{1}{2}g \succsim f$.*

Axiom 6 is a refinement of the simple diversification axiom in Siniscalchi (2009) which restricts the axiom to complementary pairs with indifferent extremes. As such, Axiom 6 permits departures from ambiguity aversion for acts that hedge utility on the best event for an act, allowing instead for a preference for speculating on ambiguous acts with high potential payoffs, consistent with the empirical findings of Baillon and Bleichrodt (2015), Dimmock et al. (2015) and Kocher et al. (2018).

Proposition 1 *Let $|S| > 3$. Axioms 1-6 are necessary and sufficient for $\succ$ to be represented by a generalized NEO-EU functional with parameters (π, λ, ρ) satisfying $(1-\rho)(\frac{1}{2}-\lambda) \geq 0$.*

A NEO-EU representation is the special case of a generalized NEO-EU representation in Theorem 2 in which $\rho, \lambda \in [0,1]$. We refer to a representation as an *ambiguity-averse NEO-EU representation* if it is a NEO-EU representation with $\lambda \in [0, \frac{1}{2})$, which is also the range of values that produce ambiguity aversion in the Ellsberg (1961) paradox. Proposition 1 implies that if $\lambda \in [0, \frac{1}{2})$, then $\rho \in [0,1]$ and the decision maker has an ambiguity-averse NEO-EU representation. Conversely, if $\rho \in [0,1)$, then $\lambda \in [0, \frac{1}{2}]$. Due to Proposition 1, one can state a representation result as follows:[6]

Corollary 1 *Let $|S| > 3$. Then for $\lambda \in [0, \frac{1}{2})$, Axioms 1-6 hold if and only if $\succ$ admits an ambiguity-averse NEO-EU representation.*

The restriction $\lambda \in [0, \frac{1}{2})$ has empirical support from existing estimates of λ in prior experiments: Dimmock et al. (2015) estimate $\lambda = 0.44$, Baillon et al. (2018a) estimate $\lambda = 0.45$, and Abdellaoui et al. (2021) estimate $\lambda = 0.18$. Further, Ghazi et al. (2023a) estimate λ for a NEO-EU representative agent of the aggregate stock market (the market's ambiguity attitude) in a generalization of the Consumption CAPM. Across their data spanning more than thirty years of monthly observations, the estimated λ fluctuates between 0 and 0.52 and the values of $\lambda > 0.5$ are concentrated in the period of the dot-com bubble.

The next definition and axiom are used in a recent working paper by Asano et al. (2022) in the Anscombe-Aumann framework. The axiom implies ambiguity seeking behavior for acts that have their worst outcome in the same state and ambiguity aversion for acts that have their best outcome in the same state. However, the axiom seems too strong as there is not a clear justification why a decision maker is or should

[6] Note that Axioms 1 - 6 imply the existence and uniqueness of λ and ρ from the axioms. This statement of the representation result which notes a range for λ has a similar form to the statement of the main theorem in Hartmann (2023) that "For each $\alpha \in [0,1]\backslash\frac{1}{2}$" his axioms are equivalent to an α-maxmin representation.

be ambiguity seeking due to acts sharing a common worst state. Asano et al. (2022) show that when added to the Choquet expected utility axioms, this additional axiom implies a NEO-EU representation with the parameters restricted to the interval $[0, 1]$. We show a related result, that Axiom 7, when combined with Axioms 1 - 5 restricts ρ and λ to the interval $[0, 1]$. However, since Axiom 7 seems too strong, we conclude that axioms leading to a NEO-EU representation with $\lambda \in [0, \frac{1}{2})$ may be more natural.

Definition 4 *Two acts $f, g \in \mathcal{F}$ are called "co-maximal" if $\overline{S}(f) = \overline{S}(g)$. Analogously, two acts $f, g \in \mathcal{F}$ are called "co-minimal" if $\underline{S}(f) = \underline{S}(g)$.*

Hence, two acts are co-maximal or co-minimal if they have the same maximum or minimum states, respectively.

Axiom 7 *(Uncertainty aversion or tolerance depending on co-extremality) For every $f, g \in \mathcal{F}$ such that $f \sim g$, f and g co-maximal implies $\frac{1}{2}f + \frac{1}{2}g \succsim f$; and f and g co-minimal implies $f \succsim \frac{1}{2}f + \frac{1}{2}g$.*

If Axiom 7 holds, a generalized NEO-EU functional is both a convex and a concave function across mixtures of two acts, ruling out parameter values outside $[0, 1]$.

Proposition 2 *Let $|S| > 3$. Axioms 1-5 and 7 are necessary and sufficient for $\succ$ to be represented by a generalized NEO-EU functional with parameter vector (π, λ, ρ) such that $\lambda, \rho \in [0, 1]$.*

The proofs of Proposition 1 and 2 are in the Appendix.

3 VC Dimension and Theory Falsifiability

Basu and Echenique (2020) study the degree of falsifiability of three central theories of choice under ambiguity: subjective expected utility (SEU), Choquet expected utility (CEU), and max-min expected utility (MEU) which is the special case of the α-maxmin multiple priors model in which $\alpha = 0$. They observe, "A theory is easy to falsify if relatively small data sets are enough to guarantee that the theory can be falsified: the Vapnik-Chervonenkis (VC) dimension of a theory is the largest sample size for which the theory is 'never falsifiable'." They consider a setting in which a strategic proponent of a theory presents experimental evidence in favor of the theory to a skeptical user of theories. They determine the smallest sample size needed for SEU, CEU, and MEU to be falsifiable, even when the experimenter can selectively design the experiment to produce favorable conditions for the theory.

As we characterize generalized NEO-EU preferences in Section 2, we ask whether our proposed axioms produce a theory that can be falsified in principle. The results of Basu and Echenique (2020) seem to indicate more broadly that axiomatic ambiguity models are difficult or impossible to falsify in practice, or at least are substantially more difficult to falsify than SEU. Here, we compute the VC dimension of generalized NEO-EU preference theories and show that their VC dimension increases linearly with the number of states, similar to SEU. Consequently, an indirect implication of the co-extreme independence axiom is that it delivers a preference representation that can be tested convincingly and efficiently with small data sets.

We denote by $\mathcal{H}$ the set of all preference relations defined on $\mathcal{F} \times \mathcal{F}$. In particular, we denote by $\mathcal{E} \subset \mathcal{H}$ the set of EU preferences, that is, $\succ \in \mathcal{E}$ if there exists $\pi \in \Delta(S)$ such that $f \succ g$ if and only if $\mu(f;\pi) > \mu(g;\pi)$. We denote by $\mathcal{N} \subset \mathcal{H}$ the set of standard NEO-EU preferences, that is, $\succ \in \mathcal{N}$ if there exists $(\pi, \lambda, \rho) \in \Delta(S) \times [0,1] \times [0,1]$ such that $f \succ g$ if and only if $V(f) > V(g)$. We denote by $\mathcal{N}_0 \subset \mathcal{H}$ the set of Hurwicz preferences, that is, $\succ \in \mathcal{N}_0$ if there exists $\lambda \in [0,1]$ such that $f \succ g$ if and only if $\nu(f) > \nu(g)$, where $\nu(f) = \nu(f;\lambda) := \lambda \max_{s \in S} U(f(s)) + (1 - \lambda) \min_{s \in S} U(f(s))$. Finally, we denote by $\mathcal{N}_1 \subset \mathcal{H}$ the set of generalized NEO-EU preferences (i.e., for which λ and ρ are not restricted to the interval [0,1], but still satisfy conditions (8)-(10) from Theorem 2). We clearly have $\mathcal{E} \cup \mathcal{N}_0 \subset \mathcal{N} \subset \mathcal{N}_1$.

Given $m \in \mathbb{N}$, let $\mathcal{D}_m := (\mathcal{F} \times \mathcal{F})^m$ be the set of samples of size m, where each coordinate in a sample (vector) is taken from $\mathcal{F} \times \mathcal{F}$. A *labeled* sample is a pair (d, b), where $d \in \mathcal{D}_m$ and $b \in \{0,1\}^m$. Given a relation $\succ$ and a labeled sample $(d, b) \in \mathcal{D}_m \times \{0,1\}^m$, with $d := \left[(f_1, \hat{f}_1), \ldots, (f_m, \hat{f}_m)\right]$, we say that $\succ$ is *consistent* with d if $f_i \succ \hat{f}_i \Leftrightarrow b_i = 1$ for all $i = 1, \ldots, m$.

Definition 5 *Given a sample vector $d \in \mathcal{D}_m$ and a set of preferences $\mathcal{C} \subset \mathcal{H}$, we say that d is shattered by $\mathcal{C}$ if for all $b \in \{0,1\}^m$ there exists a relation $\succ \in \mathcal{C}$ consistent with (d, b). Furthermore, the Vapnik-Chervonenkis (VC) dimension of $\mathcal{C}$ is defined as*

$$VC(\mathcal{C}) := \max \{m \in \mathbb{N} : \textit{there exists } d \in \mathcal{D}_m \textit{ shattered by } \mathcal{C}\}, \tag{12}$$

that is, $VC(\mathcal{C})$ is the largest positive integer m for which we can find a sample m-vector *that can be shattered by a preference in $\mathcal{C}$.*

The following theorem establishes the VC dimension of EU, standard NEO-EU, generalized NEO-EU, and Hurwicz theories.[7]

Theorem 3 *The following statements hold true:*

1. $VC(\mathcal{E}) = |S| - 1$.
2. *For* $|S| > 2$, $VC(\mathcal{N}) = VC(\mathcal{N}_1) = |S| + 1$.
3. $VC(\mathcal{N}_0) = 1$.

Proof. See Appendix. ■

Given two preferences $\succ$ and $\succ'$, and a probability distribution F on $\mathcal{F} \times \mathcal{F}$ we denote by $e_F(\succ, \succ')$ the *error* between the preferences, that is, $e_F(\succ, \succ') := P_F(\succ \triangle \succ')$, where $\triangle$ is the symmetric difference operator on $\mathcal{F} \times \mathcal{F}$. Next, we adapt standard learning theory concepts and results from Blum et al. (2020); Kearns and Vazirani (1994).

Definition 6 *We say that a set of preferences $\mathcal{C}$ (theory) is "probably approximately correct (PAC) learnable" if there exists an algorithm such that for a given preference $\succ \in \mathcal{C}$ any probability distribution F, and for any $\delta, \epsilon \in (0, 1)$, with probability at least $1 - \delta$, the algorithm, with input a random sample d (generated by using F) labeled by b consistent with $\succ$, outputs a relation $\succ' \in \mathcal{C}$ consistent with (d, b) and satisfying $e_F(\succ, \succ') < \epsilon$.*

Using Theorem 3 and Theorem 3.3 from Kearns and Vazirani (1994) (or Corollary 5.17 from Blum et al. (2020)), we obtain the following result that shows that all the preference sets considered in this section (EU, NEO-EU, generalized NEO-EU, and Hurwicz) are PAC learnable.

Corollary 2 *For a given set of preferences $\mathcal{C}$, let $\mathcal{A}$ be any algorithm that takes input a labeled sample $(d, b) \in \mathcal{D}_m \times \{0, 1\}^m$ with labels consistent with a relation in $\mathcal{C}$, and*

[7] We derive the VC dimension of EU preferences since we identified an error in the proof of the VC dimension for EU in Basu and Echenique (2020). Correcting the error, we find that the VC dimension of EU is $|S| - 1$, not $|S|$ as reported in Basu and Echenique (2020).

outputs a relation in $\mathcal{C}$ that is consistent with (d, b). Then, $\mathcal{A}$ is a PAC learning algorithm for $\mathcal{C}$ as long as the sample size m satisfies

$$m \geq \frac{c}{\epsilon}\left(VC(\mathcal{C})\log_2\frac{1}{\epsilon} + \log_2\frac{1}{\delta}\right), \tag{13}$$

for some constant $c > 0$.

The proof of Theorem 3 suggests an efficient method to determine whether there exists a standard NEO-EU relation that separates the points of a sample m-vector and, if this relation exists, how to determine the parameter vector $(\pi, \lambda, \rho) \in \Delta(S) \times [0, 1] \times [0, 1]$ of such a relation. This is summarized in the following result.

Proposition 3 *Given a sample $d := \left[(f_1, \hat{f}_1), \ldots, (f_m, \hat{f}_m)\right] \in \mathcal{D}_m$ and a nonzero label vector $b \in \{0, 1\}^m$ consistent with an agent's preferences among the acts in d, let $v^i_s := U(f_i(s))$ and $\hat{v}^i_s := U(\hat{f}_i(s))$ for all $s \in S$ and $i = 1, \ldots, m$; $A := [a^1, \ldots, a^m]$ be the matrix whose columns are $a^i := v^i - \hat{v}^i$; $g := [\max(v^i) - \max(\hat{v}^i)]_{i=1,\ldots,m}$; $h := [\min(v^i) - \min(\hat{v}^i)]_{i=1,\ldots,m}$; $I := \{i : b_i = 1\}$; and $J := \{i : b_i = 0\}$. Consider the following linear optimization problem:*

$$\begin{array}{lll} (P) & \max & \sigma, \\ & s.t. & \\ & & A_I^T w + \alpha g_I + \beta h_I \geq \sigma e_I, \\ & & A_J^T w + \alpha g_J + \beta h_J \leq 0, \\ & & e^T w + \alpha + \beta = 1, \\ & & \alpha, \beta, \sigma, w \geq 0. \end{array}$$

Then, there exists a relation $\succ$ in $\mathcal{N}$ consistent with the labeled sample (d, b) if and only if problem (P) has a solution $(\alpha, \beta, \sigma, w)$ with $\sigma > 0$. In that case $(\pi, \lambda, \rho) \in \Delta(S) \times [0, 1] \times [0, 1]$ defined as follows determines a relation $\succ$ in $\mathcal{N}$ consistent with b:

$$\begin{aligned} \rho &:= e^T w, \\ \lambda &:= \begin{cases} \frac{\alpha}{\alpha+\beta} & \text{if } 0 < \rho < 1, \\ \alpha & \text{if } \rho = 0, \\ 1 & \text{if } \rho = 1, \end{cases} \end{aligned}$$

$$\pi \ := \begin{cases} \frac{w}{\rho} & \textit{if } 0 < \rho < 1, \\ \frac{e}{|S|} & \textit{if } \rho = 0, \\ w & \textit{if } \rho = 1. \end{cases}$$

If (P) is infeasible or has an optimal solution with $\sigma = 0$, then (d, b) is proof that the standard NEO-EU theory is false for this agent. If b is a zero vector (so that I is empty), then there exists a relation $\succ$ in $\mathcal{N}$ consistent with the sample if and only if the system $A^T w + \alpha g + \beta h \leq 0$, $e^T w + \alpha + \beta = 1$, and $\alpha, \beta, w \geq 0$ is feasible.

Proof. See Appendix. ∎

Provided we can efficiently compute utility $U(f(s))$ of an act f in $\mathcal{F}$ for each $s \in S$, using a polynomial-time linear programming algorithm, and using a sample size m satisfying bound (13), this algorithm is an efficient (i.e., with polynomial running time) PAC learning algorithm for the class $\mathcal{N}$ of standard NEO-EU preferences. In the case $b \neq 0$, there is a linear program analogous to (P) to determine if an EU theory is consistent with a data sample, namely, $\max_{\sigma,w}\{\sigma : A_I^T w \geq \sigma e_I, A_J^T w \leq 0, e^T w = 1, \sigma, w \geq 0\}$. Similar to Proposition 3, there is an EU preference consistent with the data if and only if there is an optimal solution to this problem with $\sigma > 0$. Since identifying an EU preference consistent with the data and the problem (P) of identifying a NEO-EU preference consistent with the data can each be represented via linear programs (which differ only in two columns), the two problems have essentially the same computational complexity. Similar remarks apply to the case $b = 0$.

4 Conclusion

Often in applications, one seeks a simple parametric model that can be justified on axiomatic grounds. However, there is typically a tradeoff as ambiguity models based on normatively appealing axioms such as certainty independence, co-monotonic independence, or complementary independence characterize very general classes of preferences and provide little guidance on which specific functional form to adopt. Our main representation theorem (Theorem 2) demonstrates that the axiom of co-extreme independence, when combined with the standard invariant bi-separable axioms (non-trivial weak order, continuity, monotonicity, and certainty independence) delivers a precise parametric representation of preferences (the generalized NEO-EU representation) in the standard Anscombe-Aumann framework in which the existence and uniqueness of the parameters and the subjective probability distribution are determined from the axioms. We also show that a refinement of the preference for hedging axiom from Siniscalchi (2009) helps to provide the first characterization of an ambiguity-averse NEO-EU representation, which can be stated in a form similar to the statement of the main theorem characterizing α-maxmin preferences in Hartmann (2023).

We brought into our analysis a criterion for evaluating ambiguity models introduced by Basu and Echenique (2020): the degree to which ambiguity models are falsifiable. We find that although generalized NEO-EU and NEO-EU preferences weaken the standard SEU independence axiom and so admit a more general class of preferences, they each have a VC dimension that increases linearly with the number of states, making them essentially no more difficult to falsify than SEU.

By providing an axiomatic foundation for generalized NEO-EU preferences and ambiguity-averse NEO-EU preferences in the standard Anscombe-Aumann framework and determining the VC dimension of NEO-EU and the recoverability of its parameters, our approach helps clarify the foundations and properties of the NEO-EU model. To the extent the axioms are reasonable, the results presented here further justify its broad use in economic applications.

References

Abdellaoui, Mohammed, Philippe Colo, Brian Hill. 2021. Eliciting multiple prior beliefs. *HEC Paris Research Paper No. ECO/SCD-2021-1426* .

Anscombe, Francis J, Robert J Aumann. 1963. A definition of subjective probability. *Annals of mathematical statistics* **34**(1) 199–205.

Asano, Takao, Hiroyuki Kojima, Kaname Miyagishima. 2022. A simple axiomatization of neo-additive choquet expected utility theory on a finite state space. *KIER Discussion Paper* **1080** 1–31.

Baillon, Aurélien, Han Bleichrodt. 2015. Testing ambiguity models through the measurement of probabilities for gains and losses. *American Economic Journal: Microeconomics* **7**(2) 77–100.

Baillon, Aurélien, Han Bleichrodt, Umut Keskin, Olivier l'Haridon, Chen Li. 2018a. The effect of learning on ambiguity attitudes. *Management Science* **64**(5) 2181–2198.

Baillon, Aurélien, Zhenxing Huang, Asli Selim, Peter P Wakker. 2018b. Measuring ambiguity attitudes for all (natural) events. *Econometrica* 1839–1858.

Basu, Pathikrit, Federico Echenique. 2020. On the falsifiability and learnability of decision theories. *Theoretical Economics* **15**(4) 1279–1305.

Bazaraa, M.S., H.D. Sherali, C.M. Shetty. 2006. *Nonlinear Programming*. 3rd ed. John Wiley & Sons, Inc., New York, NY.

Blum, A., J. Hopcroft, K. Ravindran. 2020. *Foundations of Data Science*. Cambridge University Press, New York, NY.

Chambers, Christopher P, Federico Echenique, Nicolas S Lambert. 2021. Recovering preferences from finite data. *Econometrica* **89**(4) 1633–1664.

Chambers, Christopher P, Federico Echenique, Nicolas S Lambert. 2023. Recovering utility. *arXiv preprint arXiv:2301.11492* .

Chambers, Robert G, Simon Grant, Ben Polak, John Quiggin. 2014. A two-parameter model of dispersion aversion. *Journal of Economic Theory* **150** 611–641.

Chateauneuf, Alain, Jürgen Eichberger, Simon Grant. 2007. Choice under uncertainty with the best and worst in mind: Neo-additive capacities. *Journal of Economic Theory* **137**(1) 538–567.

Denti, Tommaso, Luciano Pomatto. 2022. Model and predictive uncertainty: A foundation for smooth ambiguity preferences. *Econometrica* **90**(2) 551–584.

Dimmock, Stephen G, Roy Kouwenberg, Olivia S Mitchell, Kim Peijnenburg. 2015. Estimating ambiguity preferences and perceptions in multiple prior models: Evidence from the field. *Journal of Risk and Uncertainty* **51** 219–244.

Dominiak, Adam, Ani Guerdjikova. 2021. Pessimism and optimism towards new discoveries. *Theory and Decision* **90** 321–370.

Dow, James, Sérgio Ribeiro da Costa Werlang. 1992. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica: Journal of the Econometric Society* 197–204.

Eichberger, Jürgen, Simon Grant, Jean-Philippe Lefort. 2012. Generalized neo-additive capacities and updating. *International Journal of Economic Theory* **8**(3) 237–257.

Eichberger, Jürgen, David Kelsey. 1999. E-capacities and the ellsberg paradox. *Theory and decision* **46** 107–138.

Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics* **75**(4) 643–669.

Frick, Mira, Ryota Iijima, Yves Le Yaouanq. 2022. Objective rationality foundations for (dynamic) α-meu. *Journal of Economic Theory* **200** 105394.

Ghazi, Soroush, Mark Schneider, Jack Strauss. 2023a. Market ambiguity attitude restores the risk-return tradeoff. *Management Science, forthcoming* .

Ghazi, Soroush, Mark Schneider, Jack Strauss. 2023b. Momentum is still there conditional on crash aversion. *Available at SSRN 4680094* .

Ghirardato, Paolo, Fabio Maccheroni, Massimo Marinacci. 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* **118**(2) 133–173.

Gilboa, Itzhak, David Schmeidler. 1989. Maxmin expected utility with non-unique prior. *Journal of mathematical economics* **18**(2) 141–153.

Giraud, Raphaël, Lionel Thomas. 2017. Ambiguity, optimism, and pessimism in adverse selection models. *Journal of Economic Theory* **171** 64–100.

Grant, S., B. Polak. 2013. Mean-dispersion preferences and constant absolute uncertainty aversion. *Journal of Economic Theory* **148** 1361–1398.

Groneck, Max, Alexander Ludwig, Alexander Zimper. 2016. A life-cycle model with ambiguous survival beliefs. *Journal of Economic Theory* **162** 137–180.

Hartmann, Lorenz. 2023. Strength of preference over complementary pairs axiomatizes alpha-meu preferences. *Journal of Economic Theory* **213** 105719.

Hurwicz, L. 1950. Optimality criteria for decision making under ignorance. Discussion Paper 370, Cowles Commission.

Kearns, M.J., U.V. Vazirani. 1994. *An Introduction to Computational Learning Theory*. The MIT Press, Cambridge, MA.

Keynes, John Maynard. 1921. A treatise on probability .

Knight, Frank H. 1921. Risk, uncertainty and profit. *Hart, Schaffner and Marx* .

Kocher, Martin G, Amrei Marie Lahno, Stefan T Trautmann. 2018. Ambiguity aversion is not universal. *European Economic Review* **101** 268–283.

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

Kreps, D.M. 1988. *Notes on the Theory of Choice*. Westview Press, Inc., Boulder, CO.

Lockwood, Benjamin, Hunt Allcott, Dmitry Taubinsky, Afras Y Sial. 2021. What drives demand for state-run lotteries? evidence and welfare implications. Tech. rep., National Bureau of Economic Research.

Maccheroni, Fabio, Massimo Marinacci, Aldo Rustichini. 2006. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* **74**(6) 1447–1498.

Mas-Collel, A., M.D. Whinston, J.R. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York, NY.

Neilson, William S. 2010. A simplified axiomatic approach to ambiguity aversion. *Journal of Risk and uncertainty* **41** 113–124.

Nunez, M., M. Schneider. 2019. Mean-dispersion preferences with a specific dispersion function. *Journal of Mathematical Economics.* **84**(1) 195–206.

Savage, L.J. 1954. *The Foundations of Statistics*. John Wiley & Sons, Inc., New York, NY.

Schmeidler, David. 1989. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society* 571–587.

Siniscalchi, Marciano. 2009. Vector expected utility and attitudes toward variation. *Econometrica* **77**(3) 801–855.

Strzalecki, Tomasz. 2011. Axiomatic foundations of multiplier preferences. *Econometrica* **79**(1) 47–73.

Toquebeuf, Pascal. 2016. Choquet expected utility with affine capacities. *Theory and Decision* **81** 177–187.

Wakker, Peter P. 2010. *Prospect theory: For risk and ambiguity*. Cambridge university press.

Wakker, Peter P, Horst Zank. 2002. A simple preference foundation of cumulative prospect theory with power utility. *European Economic Review* **46**(7) 1253–1271.

Wald, A. 1950. *Statistical Decision Functions*. John Wiley & Sons, Inc., New York, NY.

Zank, Horst. 2007. Social welfare functions with a reference income. *Social Choice and Welfare* **28**(4) 609–636.

Zimper, Alexander. 2012. Asset pricing in a lucas fruit-tree economy with the best and worst in mind. *Journal of Economic Dynamics and Control* **36**(4) 610–628.

Appendix: Proofs of Main Results

State-Utility Space

Throughout this section we assume that Axioms 1 through 4 hold. We also use the utility function u from Theorem 1 normalized as indicated in (2). We denote by $\hat{u}$ the $|X|$-dimensional vector whose coordinates are $\hat{u}_x := u(x)$, for all $x \in X$. Notice that for each act $f \in \mathcal{F}$ there is a corresponding $|S| \times |X|$-matrix $M(f)$ defined as $M(f)_{sx} = f_s(x)$. We denote by e the all-ones vector whose dimension is determined by the context in which is used. We denote by e_s the vector in $\mathbb{R}^{|S|}$ whose s-th coordinate is 1 and all other coordinates are 0. Hence, we have $M(f)e = e$ for all f. Moreover, any nonnegative $|S| \times |X|$-matrix M such that $Me = e$ corresponds to an act in $\mathcal{F}$, namely the act f defined as $f_s(x) := M_{sx}$. Therefore, the set $\mathcal{F}$ can be identified with the set of all nonnegative $|S| \times |X|$-matrices M satisfying $Me = e$. Also notice that $M(\alpha f + (1-\alpha)g) = \alpha M(f) + (1-\alpha)M(g)$ for all $\alpha \in [0,1]$.

Let $\mathcal{G} := [0,1]^{|S|}$. We call $\mathcal{G}$ the *state-utility space*. In particular, given $v \in \mathcal{G}$, we can define $f(s) := v_s \overline{x} + (1-v_s)\underline{x}$ for each $s \in S$. Using normalization (2), it follows that $U(f(s)) = v_s$ for each s. Hence, for each $v \in \mathcal{G}$ there exists $f \in \mathcal{F}$ such that $M(f)\hat{u} = v$. Following Grant and Polak (2013), on $\mathcal{G}$ we define a binary relation $\succ^u$ as follows: $v \succ^u v'$ if there exist acts f and f' in $\mathcal{F}$ such that $v = M(f)\hat{u}$, $v' = M(f')\hat{u}$, and $f \succ f'$, for all $v, v' \in \mathcal{G}$. The binary relation $\succ^u$ is well-defined and satisfies equivalent axioms to Axioms 1-4 on $\mathcal{G}$. The following is a standard result for preferences satisfying those axioms (Grant and Polak, 2013).

Proposition 4 *For each $v \in \mathcal{G}$ there exists a unique $\alpha_v \in [0,1]$ such that $v \sim^u \alpha_v e$.*

Based on Proposition 4, we define function $\varphi : \mathcal{G} \to [0,1]$ as $\varphi(v) = \alpha_v$, where α_v is the scalar indicated in the proposition. Hence, $v \sim^u \varphi(v)e$ for all $v \in \mathcal{G}$. By definition, we readily obtain $0 \leq \varphi(v) \leq 1$ for all v, $\varphi(0) = 0$, and $\varphi(e) = 1$. The next result states well-known properties of φ whose proofs are standard in the literature (see Mas-Collel et al., 1995; Grant and Polak, 2013; Nunez and Schneider, 2019).

Lemma 1 *The following properties of φ hold:*

1. *$v \succ^u v'$ if and only if $\varphi(v) > \varphi(v')$, for all $v, v' \in \mathcal{G}$;*

2. *$\varphi(\alpha v + (1-\alpha)\beta e) = \alpha\varphi(v) + (1-\alpha)\beta$ for all $v \in \mathcal{G}$ and $\alpha, \beta \in [0,1]$;*

3. *$\varphi(\alpha v) = \alpha\varphi(v)$ for all $v \in \mathcal{G}$ and $\alpha \in [0,1]$;*

4. *$\varphi(v + \delta e) = \varphi(v) + \delta$ for all $v \in \mathcal{G}$ and $\delta \in \mathbb{R}$ such that $v + \delta e \in \mathcal{G}$;*

5. *if $v \sim^u v'$ and $\delta \in \mathbb{R}$ is such that $v + \delta e, v' + \delta e \in \mathcal{G}$, then $v + \delta e \sim^u v' + \delta e$.*

6. *φ is a continuous function on $\mathcal{G}$.*

Proof of Theorem 2

Throughout this section we assume that Axioms 1 through 5 hold. We first extend the definition of co-extremality to state-utility vectors analogously to Definitions 1 and 2. Hence, for $v \in \mathcal{G}$ let $\overline{S}(v) := \{s : v_s \geq v_{s'} \text{ for all } s' \in S\}$, $\underline{S}(v) :=$

$\{s : v_s \leq v_{s'} \text{ for all } s' \in S\}$, and $R(v) := S \setminus \left(\overline{S}(v) \cup \underline{S}(v)\right)$. We say that two vectors $v, \hat{v}$ are *co-extreme* and write $v \equiv \hat{v}$, if $\overline{S}(v) = \overline{S}(\hat{v})$ and $\underline{S}(v) = \underline{S}(\hat{v})$. The maximum and minimum of a vector v are defined as usual: $\max(v) := \max\{v_s : s \in S\}$, and $\min(v) := \min\{v_s : s \in S\}$.

Observe that if v is not proportional to the e vector, then $\overline{S}(v)$ and $\underline{S}(v)$ are nonempty nonoverlapping sets. Moreover, in that case, $v_s = \max(v)$ for all $s \in \overline{S}(v)$ and $v_s = \min(v)$ for all $s \in \underline{S}(v)$. If v is proportional to e, i.e., all its entries are the same, then $\overline{S}(v) = \underline{S}(v) = S$. Also notice that it is easy to show that if $f, \bar{f} \in \mathcal{F}$ are such that $v = M(f)\hat{u}$ and $\hat{v} = M(\bar{f})\hat{u}$, then $v, \hat{v}$ are co-extreme if and only if $f, \bar{f}$ are co-extreme. The relation $\equiv$ on $\mathcal{G}$ is an equivalence relation. Therefore, $\mathcal{G}$ can be partitioned into a finite collection of equivalence classes, each equivalent class is a subset of $\mathcal{G}$. As usual, we denote by $\mathcal{G}/\equiv$ the set of equivalent classes, and call them "co-extreme" to emphasize that they are derived from the $\equiv$ relation. In particular, we denote by $\mathcal{G}_0$ the class of vectors v such that $\overline{S}(v) = \underline{S}(v) = S$, which is the same as the set of all vectors in $\mathcal{G}$ that are proportional to e. If $|S| \geq 3$, for every state $s \in S$ we can always have a vector v such that v_s is neither the maximum nor minimum of v, i.e., $s \in R(v)$. Therefore, the number of equivalence classes other than $\mathcal{G}_0$ is the same as the number of ways that we can partition S into two or three nonempty sets, i.e., $K := 3^{|S|} - 2^{|S|+1} + 1$. We denote by $\mathcal{G}_k$, $k = 1, \ldots, K$, the classes in $\mathcal{G}$ other than $\mathcal{G}_0$. Since $v \equiv \hat{v}$ implies $\overline{S}(v) = \overline{S}(\hat{v})$, we denote by $\overline{S}_k$ the common set of maximum states across all vectors $v \in \mathcal{G}_k$. Similarly, we denote by $\underline{S}_k$ the common set of minimum states across all vectors $v \in \mathcal{G}_k$. We say that a class $\mathcal{G}_k$ is *mono-extreme* if $|\overline{S}_k| = |\underline{S}_k| = 1$. Notice that if $v \in \mathcal{G}_k$ and $\mathcal{G}_k$ is mono-extreme, then there is a unique state $\overline{s}$ such that $v_{\overline{s}} = \max(v)$, a unique state $\underline{s}$ such that $v_{\underline{s}} = \min(v)$, and $v_{\overline{s}} > v_s > v_{\underline{s}}$ for all $s \in R_k$. Also, it is easy to see that the number of mono-extreme classes in $\mathcal{G}/\equiv$ is $|S|\,(|S| - 1)$. The next result follows from the maximum and minimum definitions.

Proposition 5 *Let $v, \hat{v} \in \mathcal{G}$ and $\alpha \in [0, 1]$. If $\overline{S}(v) \cap \overline{S}(\hat{v}) \neq \emptyset$, then $\max\left(\alpha v + (1-\alpha)\hat{v}\right) = \alpha \max(v) + (1-\alpha)\max(\hat{v})$. If $\underline{S}(v) \cap \underline{S}(\hat{v}) \neq \emptyset$, then $\min\left(\alpha v + (1-\alpha)\hat{v}\right) = \alpha \min(v) + (1-\alpha)\min(\hat{v})$. If $v \equiv \hat{v}$, then $\alpha v + (1-\alpha)\hat{v}$ is co-extreme with v and $\hat{v}$.*

Proposition 5 implies that each class $\mathcal{G}_k$ is a mixture space as defined in Kreps (1988). Moreover, since $\succ^u$ satisfies the preference and continuity axioms in $\mathcal{G}_k$, and Axiom 5 implies that the independence axiom is also satisfied in $\mathcal{G}_k$, from the mixture space theorem (see Theorem 5.11 in Kreps, 1988), we obtain linearity of φ across $\mathcal{G}_k$:

Proposition 6 *For all $v, \hat{v} \in \mathcal{G}_k$, and $\alpha \in [0,1]$, we have $\varphi(\alpha v + (1-\alpha)\hat{v}) = \alpha\varphi(v) + (1-\alpha)\varphi(\hat{v})$.*

The next result shows that for each class there exists a linear preference functional that represents an agent's choices in that class. If a class is mono-extreme, the vector parameter associated with the class representation is uniquely determined, but for a class not mono-extreme there could be more than one vector parameter that yields the same representation. On the other hand, all class representations reduce to a representation of a mono-extreme class, that is, in reality there are at most $|S|\,(|S|-1)$ different vector parameters that represent choices across all classes.

Lemma 2 *Suppose that the relation $\succ$ on $\mathcal{F}$ satisfies Axioms 1 through 5. Then, for $k \geq 1$ there exists $w^{(k)} \in \Delta(S)$ such that*

$$\varphi(v) = {w^{(k)}}^T v, \tag{14}$$

for all $v \in \mathcal{G}_k$. The coordinates $w_s^{(k)}$ for $s \in R_k$, the term $\sum_{s \in \overline{S}_k} w_s^{(k)}$, and the term $\sum_{s \in \underline{S}_k} w_s^{(k)}$ are uniquely determined by the class. In particular, the vector $w^{(k)}$ associated with a mono-extreme class is unique. The normalized utility function u from (2) is common to all the representations across the classes. Furthermore, if $\mathcal{G}_j$ and $\mathcal{G}_k$ are co-extreme classes such that $\overline{S}_j \subset \overline{S}_k$, and $\underline{S}_j \subset \underline{S}_k$, then ${w^{(k)}}^T v = {w^{(j)}}^T v$, for all $v \in \mathcal{G}_k$.

Proof. We will show that there is a linear representation in state-utility space for each class $\mathcal{G}_k$. In doing so, we will assume that R_k is not empty because, except for a few minor details, the proof is almost identical when R_k is empty. First, notice that $\mathcal{G}_k$ is a convex set. This is because if $v, \hat{v} \in \mathcal{G}_k$ and $\alpha \in [0,1]$, then by Proposition 5, $\alpha v + (1-\alpha)\hat{v} \equiv v \equiv \hat{v}$. Therefore, $\alpha v + (1-\alpha)\hat{v} \in \mathcal{G}_k$, so that $\mathcal{G}_k$ is convex.

Next, without loss of generality, we assume that each $v \in \mathcal{G}_k$ is of the form $v = \left(v_{\overline{S}_k}, v_{\underline{S}_k}, v_{R_k}\right)$, that is, the first $|\overline{S}_k|$ coordinates are in $\overline{S}_k$, the next $|\underline{S}_k|$ coordinates are in $\underline{S}_k$, and the last $|R_k|$ coordinates are in R_k.

Let $\phi : \mathcal{G}_k \to \mathbb{R}^{2+|R_k|}$ be the mapping defined as $\phi(v) = (\max(v), \min(v), v_{R_k})$ and $\mathcal{Y}$ the image of ϕ, that is, $\mathcal{Y} := \phi(\mathcal{G}_k)$. It is easy to show that ϕ is bijective mapping between $\mathcal{G}_k$ and $\mathcal{Y}$, so that $\phi^{-1} : \mathcal{Y} \to \mathcal{G}_k$ exists. Moreover, both ϕ and ϕ^{-1}

are continuous functions because the max and min functions are continuous. The convexity of $\mathcal{G}_k$ and Proposition 6 imply that $\mathcal{Y}$ is also a convex set. Furthermore, we have $\phi(\alpha v + (1-\alpha)v') = \alpha\phi(v) + (1-\alpha)\phi(v')$, for all $v, v' \in \mathcal{G}_k$, and $\alpha \in [0,1]$. Since ϕ is bijective, it follows that $\phi^{-1}(\alpha y + (1-\alpha)y') = \alpha\phi^{-1}(y) + (1-\alpha)\phi^{-1}(y')$, for all $y, y' \in \mathcal{Y}$, and $\alpha \in [0,1]$. By Proposition 6, it follows that the real-valued function $\psi := \varphi \circ \phi^{-1}$ is continuous on $\mathcal{Y}$ and satisfies $\psi(\alpha y + (1-\alpha)y') = \alpha\psi(y) + (1-\alpha)\psi(y')$, for all $y, y' \in \mathcal{Y}$, and $\alpha \in [0,1]$.

Next, we show that $\mathcal{Y}$ has a nonempty interior. Let $y' := (3/4, 1/4, \overbrace{1/2, \ldots, 1/2}^{|R_k| \text{ times}})$. Let v' be defined as

$$v'_s := \begin{cases} 3/4 & \text{if } s \in \overline{S}_k, \\ 1/4 & \text{if } s \in \underline{S}_k, \\ 1/2 & \text{if } s \in R_k. \end{cases}$$

Then clearly $v' \in \mathcal{G}_k$, $y'_s = v'_s = 1/2$ for $s \in R_k$, $y'_1 = \max(v') = 3/4$, $y'_2 = \min(v') = 1/4$, and $y' \in \mathcal{Y}$. Moreover, for $\epsilon \in (0, 1/8)$, we have $y' + z \in \mathcal{U}$ for all z such that $|z_s| < \epsilon$ for all $s \in S$. Hence, $y' \in \operatorname{int}(\mathcal{Y})$, so that $\operatorname{int}(\mathcal{Y})$ is nonempty.

Therefore (see Bazaraa et al., 2006), there exist unique $b := (\overline{b}, \underline{b}, b_{R_k}) \in \mathbb{R}^{2+|R_k|}$ and $\gamma \in \mathbb{R}$ such that $\psi(y) = b^T y + \gamma$ for all $y \in \mathcal{Y}$. Given $v \in \mathcal{G}_k$, it follows that $\varphi(v) = \varphi(\phi^{-1}(\phi(v))) = \psi(\phi(v)) = \overline{b}\max(v) + \underline{b}\min(v) + \sum_{s \in R_k} b_s v_s + \gamma$.

For $\epsilon \in (0,1)$, let $v^{(\epsilon)}$ be defined as

$$v^{(\epsilon)}_s := \begin{cases} 1 & \text{if } s \in \overline{S}_k, \\ 1-\epsilon & \text{if } s \in \underline{S}_k, \\ 1-\epsilon/2 & \text{if } s \in R_k. \end{cases}$$

Hence, $v^{(\epsilon)} \in \mathcal{G}_k$ for all ϵ, and $v^{(\epsilon)} \to e$ as $\epsilon \to 0^+$. By the continuity of φ, we get $\overline{b} + \underline{b} + \sum_{s \in R_k} b_s + \gamma = \lim_{\epsilon \to 0^+} \varphi\left(v^{(\epsilon)}\right) = \varphi(e) = 1$. Similarly, if $v^{(\epsilon)}$ is instead defined as

$$v^{(\epsilon)}_s := \begin{cases} 1/2 & \text{if } s \in \overline{S}_k, \\ 1/2-\epsilon & \text{if } s \in \underline{S}_k, \\ 1/2-\epsilon/2 & \text{if } s \in R_k, \end{cases}$$

for all $\epsilon \in (0, 1/2)$, then $v^{(\epsilon)} \in \mathcal{G}_k$ for all ϵ, and $v^{(\epsilon)} \to e/2$ as $\epsilon \to 0^+$. Thus, we obtain $\frac{1}{2}\left(\overline{b} + \underline{b} + \sum_{s \in R_k} b_s\right) + \gamma = \frac{1}{2}$. Therefore, combining these results, we obtain

$\overline{b} + \underline{b} + \sum_{s \in R_k} b_s = 1$ and $\gamma = 0$. Next, consider $v^{(\epsilon)}$ defined as

$$v_s^{(\epsilon)} := \begin{cases} 1 & \text{if } s \in \overline{S}_k, \\ 0 & \text{if } s \in \underline{S}_k, \\ \epsilon & \text{if } s \in R_k. \end{cases}$$

Then, we have $\overline{b} + \epsilon \sum_{s \in R_k} b_s = \varphi(v^{(\epsilon)}) \geq 0$ for all $\epsilon \in (0,1)$. Therefore, $\overline{b} \geq 0$. Next, consider $v^{(\epsilon)}$ defined as

$$v_s^{(\epsilon)} := \begin{cases} 1 & \text{if } s \in \overline{S}_k, \\ 0 & \text{if } s \in \underline{S}_k, \\ 1-\epsilon & \text{if } s \in R_k. \end{cases}$$

Then, we have $\overline{b} + (1-\epsilon) \sum_{s \in R_k} b_s = \varphi(v^{(\epsilon)}) \leq 1$ for all $\epsilon \in (0,1)$. Therefore, $\underline{b} = 1 - \overline{b} - \sum_{s \in R_k} b_s \geq 0$. Furthermore, for $s' \in R_k$, let $v^{(\epsilon)}$ be defined as

$$v_s^{(\epsilon)} := \begin{cases} 1 & \text{if } s \in \overline{S}_k, \\ 0 & \text{if } s \in \underline{S}_k, \\ 1/2 & \text{if } s \in R_k, s \neq s', \\ 1/2+\epsilon & \text{if } s = s', \end{cases}$$

for all $\epsilon \in [0, 1/2)$. Axiom 4 implies $\varphi(v^{(\epsilon)}) \geq \varphi(v^{(0)})$ for all ϵ, that is, $\overline{b} + 1/2 \sum_{s \in R_k} b_s + \epsilon b_{s'} \geq \overline{b} + 1/2 \sum_{s \in R_k} b_s$, from which we obtain $\epsilon b_{s'} \geq 0$ for all ϵ, and so, $b_{s'} \geq 0$.

Summarizing, we have proven that there exists a unique vector $b := (\overline{b}, \underline{b}, b_{R_k}) \in \mathbb{R}^{2+|R_k|}$ such that $\varphi(v) = \overline{b}\max(v) + \underline{b}\min(v) + \sum_{s \in R_k} b_s v_s$, where $b \geq 0$ and $\overline{b} + \underline{b} + \sum_{s \in R_k} b_s = 1$. Therefore, if we define $w^{(k)} \in \mathbb{R}^{|S|}$ such that $w_s^{(k)} := \overline{b}/|\overline{S}_k|$ for all $s \in \overline{S}_k$, $w_s^{(k)} := \underline{b}/|\underline{S}_k|$ for all $s \in \underline{S}_k$, and the same as b in R_k, we obtain $\sum_{s \in \overline{S}} w_s^{(k)} = \overline{b}$, $\sum_{s \in \underline{S}} w_s^{(k)} = \underline{b}$, $w_s^{(k)} = b_s$ for all $s \in R_k$. Therefore, we have $\varphi(v) = {w^{(k)}}^T v$ with $w^{(k)} \in \Delta(S)$, and we obtain (14).

Finally, to prove the last statement in the theorem, let $\mathcal{G}_j$ and $\mathcal{G}_k$ be such that $\overline{S}_j \subset \overline{S}_k$, and $\underline{S}_j \subset \underline{S}_k$. If $R_j = \emptyset$, then $S = \overline{S}_j \cup \underline{S}_j \subset \overline{S}_k \cup \underline{S}_k \subset S$, which implies $\overline{S}_k \cup \underline{S}_k = S$ and so, $R_k = \emptyset$. Therefore, in this case $\overline{S}_j = \overline{S}_k$ and $\underline{S}_j = \underline{S}_k$, so that $\mathcal{G}_j = \mathcal{G}_k$ and the result is trivially true.

Hence, for the rest of the proof we assume that $R_j \neq \emptyset$. Also notice that $R_k \subset R_j$.

Given $\epsilon > 0$, consider vectors $v(\epsilon)$ and $\hat{v}$ in $\mathcal{G}$ defined as

$$v_s(\epsilon) := \begin{cases} 1 & \text{if } s \in \overline{S}_j, \\ 1-\epsilon & \text{if } s \in \overline{S}_k \setminus \overline{S}_j, \\ \delta_s & \text{if } s \in R_k, \\ \epsilon & \text{if } s \in \underline{S}_k \setminus \underline{S}_j, \\ 0 & \text{if } s \in \underline{S}_j, \end{cases} \quad \text{and } \hat{v}_s := \begin{cases} 1 & \text{if } s \in \overline{S}_k, \\ \delta_s & \text{if } s \in R_k, \\ 0 & \text{if } s \in \underline{S}_k, \end{cases}$$

where δ_s for all $s \in R_k$ are fixed, arbitrarily chosen, scalars in $(0, 1)$. Notice that $v(\epsilon) \in \mathcal{G}_j$ for all $\epsilon < 1 - \max\{\delta_s : s \in R_k\}$, and $\hat{v} \in \mathcal{G}_k$. Moreover, $v(\epsilon) \to \hat{v}$ as $\epsilon \to 0^+$. By continuity, it follows that $\varphi(v(\epsilon)) \to \varphi(\hat{v})$ as $\epsilon \to 0^+$. Since $\varphi(v(\epsilon)) = \sum_{s\in\overline{S}_j} w_s^{(j)} + (1-\epsilon)\sum_{s\in\overline{S}_k\setminus\overline{S}_j} w_s^{(j)} + \sum_{s\in R_k} w_s^{(j)}\delta_s + \epsilon\sum_{s\in\underline{S}_k\setminus\underline{S}_j} w_s^{(j)}$, and $\varphi(\hat{v}) = \sum_{s\in\overline{S}_k} w_s^{(k)} + \sum_{s\in R_k} w_s^{(k)}\delta_s$, by letting $\epsilon \to 0^+$ we obtain

$$\sum_{s\in\overline{S}_k} w_s^{(j)} + \sum_{s\in R_k} w_s^{(j)}\delta_s = \sum_{s\in\overline{S}_k} w_s^{(k)} + \sum_{s\in R_k} w_s^{(k)}\delta_s. \tag{15}$$

Since the δ_s are arbitrary in $(0, 1)$, it follows from (15) that $\sum_{s\in\overline{S}_k} w_s^{(j)} = \sum_{s\in\overline{S}_k} w_s^{(k)}$. Using this identity to simplify (15), we obtain $\sum_{s\in R_k} w_s^{(j)}\delta_s = \sum_{s\in R_k} w_s^{(k)}\delta_s$. Once again because the δ_s are arbitrary in $(0, 1)$, we obtain $w_s^{(j)} = w_s^{(k)}$ for all $s \in R_k$. Hence, by using probability complements, we obtain $\sum_{s\in\underline{S}_k} w_s^{(j)} = \sum_{s\in\underline{S}_k} w_s^{(k)}$. Therefore, for $v \in \mathcal{G}_k$ we have ${w^{(k)}}^T v = \max(v)\sum_{s\in\overline{S}_k} w_s^{(k)} + \min(v)\sum_{s\in\underline{S}_k} w_s^{(k)} + \sum_{s\in R_k} w_s^{(k)} v_s = \max(v)\sum_{s\in\overline{S}_k} w_s^{(j)} + \min(v)\sum_{s\in\underline{S}_k} w_s^{(j)} + \sum_{s\in R_k} w_s^{(j)} v_s = {w^{(j)}}^T v$, and the result follows. $\blacksquare$

Since a vector parameter associated with a class reduces to a vector parameter associated with a mono-extreme class, we only consider mono-extreme classes in what follows. The following result provides a necessary and sufficient condition for the existence of a NEO-EU representation across mono-extreme classes.

Lemma 3 *An agent has a NEO-EU preference functional V with parameter vector $(\pi, \lambda, \rho) \in \Delta(S) \times \mathbb{R} \times \mathbb{R}$ satisfying conditions (8)-(10) if and only if for each mono-extreme class $\mathcal{G}_k$ the parameters satisfy the following system of equations:*

$$\begin{aligned} \rho\pi_{\overline{s}_k} + (1-\rho)\lambda &= w_{\overline{s}_k}^{(k)}, & (16) \\ \rho\pi_{\underline{s}_k} + (1-\rho)(1-\lambda) &= w_{\underline{s}_k}^{(k)}, & (17) \end{aligned}$$

$$\rho\pi_s \;=\; w_s^{(k)}, \forall s \in R_k, \tag{18}$$

where $w^{(k)}$ is the EU vector associated with $\mathcal{G}_k$, $\overline{S}_k = \{\overline{s}_k\}$, and $\underline{S}_k = \{\underline{s}_k\}$.

Proof. Suppose that the agent has a NEO-EU preference functional V with parameter vector (π, λ, ρ) satisfying conditions (8)-(10). Given a mono-extreme class $\mathcal{G}_k$ with $\overline{S}_k = \{\overline{s}_k\}$ and $\underline{S}_k = \{\underline{s}_k\}$, notice that the vector with coordinates $\rho\pi_{\overline{s}_k} + (1-\rho)\lambda$, $\rho\pi_{\underline{s}_k} + (1-\rho)(1-\lambda)$, and $\rho\pi_s$ for all $s \in R_k$ is in $\Delta(S)$. Moreover, we have $V(f) = \rho\pi_{\overline{s}_k} v_{\overline{s}_k} + \rho\pi_{\underline{s}_k} v_{\underline{s}_k} + \rho\sum_{s\in R_k}\pi_s v_s + (1-\rho)\left(\lambda\max(v) + (1-\lambda)\min(v)\right) = \left(\rho\pi_{\overline{s}_k} + (1-\rho)\lambda\right) v_{\overline{s}_k} + \left(\rho\pi_{\underline{s}_k} + (1-\rho)(1-\lambda)\right) v_{\underline{s}_k} + \sum_{s\in R_k}\rho\pi_s v_s$, for each $f \in \mathcal{F}$ such that $v := M(f)\hat{u} \in \mathcal{G}_k$. Hence, the parameters π, λ, and ρ yield another EU representation on $\mathcal{G}_k$. Therefore, from the uniqueness of vector $w^{(k)}$ as shown in Lemma 2, it follows that (π, λ, ρ) must satisfy system (16)-(18).

Conversely, suppose that the parameter vector (π, λ, ρ) satisfies system (16)-(18) for each mono-extreme class. Given $f, g \in \mathcal{F}$, let $v = M(f)\hat{u}$ and $v' = M(g)\hat{u}$. We have $f \succ g$ if and only if $\varphi(v) > \varphi(v')$. On the other hand, by Lemma 2, there exist mono-extreme classes $\mathcal{G}_j$ and $\mathcal{G}_k$ with respective EU vectors $w^{(j)}$ and $w^{(k)}$ such that $\varphi(v) = {w^{(j)}}^T v = \rho\sum_{s\in S}\pi_s v_s + (1-\rho)\left(\lambda\max(v) + (1-\lambda)\min(v)\right)$, and $\varphi(v') = {w^{(k)}}^T v' = \rho\sum_{s\in S}\pi_s v'_s + (1-\rho)\left(\lambda\max(v') + (1-\lambda)\min(v')\right)$. Hence, if we define functional V with parameter vector (π, λ, ρ) as in (7), we have $f \succ g$ if and only if $V(f) > V(g)$, and the result follows. ■

Lemma 4 *Let $\mathcal{G}_j$ and $\mathcal{G}_k$ be two different mono-extreme classes with respective EU vectors $w^{(j)}$ and $w^{(k)} \in \Delta(S)$, the following statements hold true:*

1. *if $\overline{S}_j = \overline{S}_k = \{s\}$, then $w_s^{(j)} = w_s^{(k)}$, and $w_{\tilde{s}}^{(j)} + w_{\hat{s}}^{(j)} = w_{\tilde{s}}^{(k)} + w_{\hat{s}}^{(k)}$ for $\tilde{s} \in \underline{S}_j$ and $\hat{s} \in \underline{S}_k$;*

2. *if $\underline{S}_j = \underline{S}_k = \{s\}$, then $w_s^{(j)} = w_s^{(k)}$, and $w_{\tilde{s}}^{(j)} + w_{\hat{s}}^{(j)} = w_{\tilde{s}}^{(k)} + w_{\hat{s}}^{(k)}$ for $\tilde{s} \in \overline{S}_j$ and $\hat{s} \in \overline{S}_k$;*

3. *if $|S| > 3$, then $w_s^{(j)} = w_s^{(k)}$ for all $s \in R_j \cap R_k$.*

Proof. For the first statement, suppose that $\overline{S}_j = \overline{S}_k = \{s\}$. Let $v \in \mathcal{G}$ be defined as

$$v_{s'} = \begin{cases} 1 & \text{if } s' = s, \\ 0 & \text{if } s' \neq s, \end{cases}$$

for each $s' \in S$. Notice that $\overline{S}_j = \overline{S}_k = \overline{S}(v)$ and $\underline{S}_j, \underline{S}_k \subset \underline{S}(v)$. Hence, from Lemma 2, we obtain $w_s^{(j)} = {w^{(j)}}^T v = \varphi(v) = {w^{(k)}}^T v = w_s^{(k)}$, as stated in statement 1. Next, suppose that $\underline{S}_j = \{\tilde{s}\}$ and $\underline{S}_k = \{\hat{s}\}$. For each $s' \in R_j \cap R_k$ we choose an arbitrary $\epsilon_{s'} \in (0,1)$. Let $v \in \mathcal{G}$ be defined as

$$v_{s'} = \begin{cases} 1 & \text{if } s' = s, \\ \epsilon_s & \text{if } s' \in R_j \cap R_k, \\ 0 & \text{if } s' = \tilde{s} \text{ or } \hat{s}. \end{cases}$$

Notice that $\overline{S}_j = \overline{S}_k = \overline{S}(v)$, and $\underline{S}_j, \underline{S}_k \subset \underline{S}(v)$. Hence, from Lemma 2, we obtain $w_s^{(j)} + \sum_{s' \in R_j \setminus \{\hat{s}\}} w_{s'}^{(j)} \epsilon_{s'} = {w^{(j)}}^T v = \varphi(v) = {w^{(k)}}^T v = w_s^{(k)} + \sum_{s' \in R_k \setminus \{\tilde{s}\}} w_{s'}^{(k)} \epsilon_{s'}$. Since $w_s^{(j)} = w_s^{(k)}$, we obtain $\sum_{s' \in R_j \cap R_k} \left(w_{s'}^{(j)} - w_{s'}^{(k)} \right) \epsilon_{s'} = 0$. Because the $\epsilon_{s'}$ are arbitrarily chosen in $(0,1)$, it follows that $w_{s'}^{(j)} = w_{s'}^{(k)}$ for all $s' \in R_j \cap R_k$ in this case. Hence, since ${w^{(j)}}^T e = {w^{(k)}}^T e = 1$, we obtain $w_{\tilde{s}}^{(j)} + w_{\hat{s}}^{(j)} = w_{\tilde{s}}^{(k)} + w_{\hat{s}}^{(k)}$, and statement 1 follows.

To prove statement 2, suppose that $\underline{S}_j = \underline{S}_k = \{s\}$. Let v be defined this time as

$$v_{s'} = \begin{cases} 0 & \text{if } s' = s, \\ 1 & \text{if } s' \neq s, \end{cases}$$

for each $s' \in S$. Notice that $\underline{S}_j = \underline{S}_k = \underline{S}(v)$ and $\overline{S}_j, \overline{S}_k \subset \overline{S}(v)$. Hence, again from Lemma 2, we obtain $1 - w_s^{(j)} = {w^{(j)}}^T v = \varphi(v) = {w^{(k)}}^T v = 1 - w_s^{(k)}$, from which we obtain $w_s^{(j)} = w_s^{(k)}$. Analogously to the proof of statement 1, we can show that in this case $w_{s'}^{(j)} = w_{s'}^{(k)}$ for all $s' \in R_j \cap R_k$. Hence, once again since ${w^{(j)}}^T e = {w^{(k)}}^T e = 1$, we obtain $w_{\tilde{s}}^{(j)} + w_{\hat{s}}^{(j)} = w_{\tilde{s}}^{(k)} + w_{\hat{s}}^{(k)}$, and statement 2 follows.

Finally, for statement 3, we assume that $R_j \cap R_k$ is not empty, otherwise there is nothing to prove. We have already shown that the statement follows whenever $\overline{S}_j = \overline{S}_k$ or $\underline{S}_j = \underline{S}_k$. Hence, suppose that $\overline{S}_j \neq \overline{S}_k$ and $\underline{S}_j \neq \underline{S}_k$. Let $\overline{s}_j, \underline{s}_j, \overline{s}_k, \underline{s}_k \in S$ be such that $\overline{S}_j = \{\overline{s}_j\}, \underline{S}_j = \{\underline{s}_j\}, \overline{S}_k = \{\overline{s}_k\}, \underline{S}_k = \{\underline{s}_k\}$. Let $s \in R_j \cap R_k$. If $\overline{s}_j \neq \underline{s}_k$, then let $\mathcal{F}_l$ be the mono-extreme class with $\overline{S}_l = \{\overline{s}_j\}$ and $\underline{S}_l = \{\underline{s}_k\}$. Notice that $s \in R_j \cap R_l$ and $s \in R_k \cap R_l$. Hence, by statements 1 and 2, we must have $w_s^{(j)} = w_s^{(l)} = w_s^{(k)}$, and the result follows. Similarly, if $\underline{s}_j \neq \overline{s}_k$, then an analogous argument shows that $w_s^{(j)} = w_s^{(k)}$. Finally, if $\overline{s}_j = \underline{s}_k$ and $\underline{s}_j = \overline{s}_k$, then consider the mono-extreme classes $\mathcal{G}_l$ and $\mathcal{G}_m$ with $\overline{S}_l = \{\overline{s}_j\}$, $\underline{S}_l = \{s'\}$, $\overline{S}_m = \{\underline{s}_j\}$, and $\underline{S}_m = \{s'\}$, where s' is any other state different from $s, \overline{s}_j, \underline{s}_j$ (this is possible because

$|S| > 3$). Then, again from statements 1 and 2, we obtain $w_s^{(j)} = w_s^{(l)} = w_s^{(m)} = w_s^{(k)}$, and the result follows. ■

Proof of Theorem 2 (axioms are sufficient): First, we show that Axioms 1-5 are sufficient, so that we assume that they hold. Given $f, g \in \mathcal{F}$ and using the same notation as in the preceding paragraphs, let $v := M(f)\hat{u}$ and $\hat{v} := M(g)\hat{u}$. Then, from Lemma 1, we have $f \succ g$ if and only if $v \succ^u \hat{v}$ iff $\varphi(v) > \varphi(\hat{v})$. So that if we show that there exist parameters $\pi \in \Delta(S)$ and $\lambda, \rho \geq 0$ satisfying (8)-(10) such that

$$\varphi(v) = \rho\pi^T v + (1-\rho)\left(\lambda \max(v) + (1-\lambda)\min(v)\right) \tag{19}$$

for all $v \in \mathcal{G}$, then by taking $V(f) := \varphi(M(f)\hat{u})$ and $V(g) := \varphi(M(g)\hat{u})$, we will obtain (11) for the normalized utility function u used in this section. Moreover, it is not difficult to see that from this we obtain (11) for any other positive affine transformation of u, thus proving that the axioms are sufficient. Hence, we concentrate on proving that (19) holds.

Since $|S| > 3$, the third statement in Lemma 4 implies that the coefficients $w_s^{(k)}$ of a mono-extreme class $\mathcal{G}_k$ that correspond to states $s \in R_k$ are independent of the class. In other words, for each $s \in S$ there exists a scalar $b_s \geq 0$ such that $w_s^{(k)} = b_s$ for all mono-extreme classes $\mathcal{G}_k$ with $s \in R_k$. Based on this, we define

$$\rho := \sum_{s \in S} b_s.$$

If $\rho > 0$, then we define

$$\pi_s := \frac{b_s}{\sum_{s' \in S} b_{s'}} = \frac{b_s}{\rho},$$

for each $s \in S$; and if $\rho = 0$, then we arbitrarily pick any vector $\pi \in \Delta(S)$. Either way, we obtain $\rho \geq 0$, $\pi \in \Delta(S)$, and

$$\rho\pi_s = w_s^{(k)},$$

for each $s \in R_k$ and for each $\mathcal{G}_k$.

If $\rho = 1$, then we arbitrarily pick λ to be any number in $[0, 1]$. So that assume that $\rho \neq 1$. Given an arbitrary mono-extreme class $\mathcal{G}_k$ with $\overline{S}_k = \{\overline{s}\}$ and $\underline{S}_k = \{\underline{s}\}$, let

$$\lambda := \frac{w^{(k)}_{\overline{s}} - \rho\pi_{\overline{s}}}{1-\rho}.$$

Notice that

$$\begin{aligned}
1-\lambda &= \frac{1}{1-\rho}\left(1 - w^{(k)}_{\overline{s}} - \rho(1-\pi_{\overline{s}})\right) \\
&= \frac{1}{1-\rho}\left(w^{(k)}_{\underline{s}} + \sum_{s\in R_k} w^{(k)}_{s} - \rho\pi_{\underline{s}} - \sum_{s\in R_k}\rho\pi_s\right) \\
&= \frac{w^{(k)}_{\underline{s}} - \rho\pi_{\underline{s}}}{1-\rho}.
\end{aligned}$$

Now, let $\mathcal{G}_{k'}$ be another mono-extreme class different from $\mathcal{G}_k$ with $\overline{S}_{k'} = \{\hat{s}\}$ and $\underline{S}_{k'} = \{\check{s}\}$. From Lemma 4 statements 1 and 2, if $\overline{s} = \hat{s}$ or $\underline{s} = \check{s}$, then we obtain $w^{(k)}_{\overline{s}} = w^{(k')}_{\hat{s}}$ or $w^{(k)}_{\underline{s}} = w^{(k')}_{\check{s}}$. In either case, it follows that

$$\frac{w^{(k')}_{\hat{s}} - \rho\pi_{\hat{s}}}{1-\rho} = \frac{w^{(k)}_{\overline{s}} - \rho\pi_{\overline{s}}}{1-\rho} = \lambda.$$

If $\overline{s} \neq \hat{s}$ and $\underline{s} \neq \check{s}$, then consider the mono-extreme class $\mathcal{G}_{k''}$ with $\overline{S}_{k''} = \{\hat{s}\}$ and $\underline{S}_{k''} = \{\underline{s}\}$. Again from Lemma 4 statements 1 and 2, it follows that

$$w^{(k)}_{\overline{s}} + w^{(k)}_{\hat{s}} = w^{(k'')}_{\overline{s}} + w^{(k'')}_{\hat{s}} = w^{(k')}_{\overline{s}} + w^{(k')}_{\hat{s}}.$$

Thus,

$$\frac{w^{(k')}_{\hat{s}} - \rho\pi_{\hat{s}}}{1-\rho} = \frac{w^{(k')}_{\hat{s}} - w^{(k)}_{\hat{s}}}{1-\rho} = \frac{w^{(k)}_{\overline{s}} - w^{(k'')}_{\overline{s}}}{1-\rho} = \frac{w^{(k)}_{\overline{s}} - \rho\pi_{\overline{s}}}{1-\rho} = \lambda.$$

Therefore, we have proven that the identities

$$\rho\pi_{\overline{s}_k} + (1-\rho)\lambda = w^{(k)}_{\overline{s}_k}, \tag{20}$$

$$\rho\pi_{\underline{s}_k} + (1-\rho)\lambda = w^{(k)}_{\underline{s}_k}, \tag{21}$$

hold for any mono-extreme class $\mathcal{G}_k$ with $\overline{S}_k = \{\overline{s}_k\}$ and $\underline{S}_k = \{\underline{s}_k\}$. Hence, from Lemma 3, it follows that the agent has a NEO-EU preference functional V with parameters (π, λ, ρ) as defined above. Moreover, $\rho \geq 0$, and since for each $s \in S$ there exists a mono-extreme class with $\overline{S} = \{s\}$ and another mono-extreme class with $\underline{S} = \{s\}$, it follows from identities (20) and (21) that $\rho\pi_s + (1-\rho)\lambda \geq 0$ and $\rho\pi_s + (1-\rho)(1-\lambda) \geq 0$ for all $s \in S$.

From Theorem 1, it follows that the utility function u is nonconstant and unique up to a positive linear transformation. Next, we show that ρ is unique. For suppose that there is another NEO-EU representation $\hat{V}$ in the form (7) with parameters $(\hat{\pi}, \hat{\lambda}, \hat{\rho})$ and utility function $\hat{u}$. By normalizing $\hat{u}$, we can assume without loss of generality that $\hat{u} = u$. Hence, $\hat{u}$ generates the same state-utility space as $\mathcal{G}$. Let $\hat{\varphi}$ be the representation of $\hat{V}$ in $\mathcal{G}$, that is,

$$\hat{\varphi}(v) = \hat{\rho}\hat{\pi}^T v + (1-\hat{\rho})\left(\hat{\lambda}\max(v) + (1-\hat{\lambda})\min(v)\right),$$

for all $v \in \mathcal{G}$. Since for any $v \in \mathcal{G}$ we have $v \sim^u \varphi(v)e$, it follows that

$$\hat{\rho}\hat{\pi}^T v + (1-\hat{\rho})\left(\hat{\lambda}\max(v) + (1-\hat{\lambda})\min(v)\right) = \hat{\varphi}(v) = \hat{\varphi}(\varphi(v)e) = \varphi(v)$$

$$= \rho\pi^T v + (1-\rho)\left(\lambda\max(v) + (1-\lambda)\min(v)\right),$$

for all $v \in \mathcal{G}$. By considering all vectors in $\mathcal{G}$ that have exactly one coordinate equal to 1 and the other coordinates equal to 0, it follows that,

$$\rho\pi_s + (1-\rho)\lambda = \hat{\rho}\hat{\pi}_s + (1-\hat{\rho})\hat{\lambda}, \tag{22}$$

for all $s \in S$. Similarly, by considering all vectors in $\mathcal{G}$ that have exactly one coordinate equal to 0 and the other coordinates equal to 1, it follows that,

$$\rho(1-\pi_s) + (1-\rho)\lambda = \hat{\rho}(1-\hat{\pi}_s) + (1-\hat{\rho})\hat{\lambda},$$

for all $s \in S$. Adding these identities for all s, we get

$$\begin{aligned} \rho + |S|(1-\rho)\lambda &= \hat{\rho} + |S|(1-\hat{\rho})\hat{\lambda}, \\ \rho(|S|-1) + |S|(1-\rho)\lambda &= \hat{\rho}(|S|-1) + |S|(1-\hat{\rho})\hat{\lambda}. \end{aligned}$$

Thus,

$$\rho - \hat{\rho} = |S|(1-\hat{\rho})\hat{\lambda} - |S|(1-\rho)\lambda = (|S|-1)(\rho - \hat{\rho}).$$

Since we are assuming that $|S| > 3$, we obtain $\rho = \hat{\rho}$, and so the ρ is unique. Using that $\rho = \hat{\rho}$, we also obtain

$$|S|(1-\rho)(\lambda - \hat{\lambda}) = 0$$

Therefore, if $\rho \neq 1$, we obtain $\lambda = \hat{\lambda}$, so that λ is unique. We also obtain from (22) that, if $\rho > 0$, $\pi_s = \hat{\pi}_s$ for all $s \in S$, and so, the π is unique. ■

Proof of Theorem 2 (axioms are necessary): Suppose that there exists a NEO-EU preference functional V with parameters (π, λ, ρ) satisfying conditions (8) to (10), with nonconstant utility function u, and such that $f \succ g$ if and only if $V(f) > V(g)$, for all $f, g \in \mathcal{F}$. Since u is nonconstant, it is easy to verify that V defines a nontrivial preference relation on $\mathcal{F}$, so that Axiom 1 holds. Without loss of generality, we assume that u has been normalized such that $u(\overline{x}) = 1$ and $u(\underline{x}) = 0$. Hence, we can define state-utility space $\mathcal{G}$ in relation to u as usual, i.e., $v \in \mathcal{G}$ if and only if there exists $f \in \mathcal{F}$ such that $M(f)\hat{u} = v$. Let $\varphi : \mathcal{G} \to \mathbb{R}$ be defined as $\varphi(v) := \rho\pi^T v + (1-\rho)\left(\lambda \max(v) + (1-\lambda)\min(v)\right)$, for all $v \in \mathcal{G}$, so that $V(f) = \varphi(M(f)\hat{u})$, for all $f \in \mathcal{F}$. Thus, Axiom 2 is an easy consequence from φ being a continuous function on $\mathcal{G}$. Axiom 3 follows from noticing that $\varphi(\alpha v + (1-\alpha)\beta e) = \alpha\varphi(v) + (1-\alpha)\beta$, for all $\alpha, \beta \in [0, 1]$.

Next, consider the following optimization problem:

$$z := \max\{\varphi(\hat{v}) - \varphi(v) : v \geq \hat{v}, v, \hat{v} \in \mathcal{G}\}. \tag{23}$$

Notice that by taking $v = \hat{v}$, we can obtain a feasible solution with objective equal to 0, so that $z \geq 0$. It follows that Axiom 4 holds if and only if $z = 0$. To show that $z = 0$, notice that problem (23) can be solved by solving $O(|S|^4)$ linear programs of the form

$$\begin{aligned}
z_{ijkl} := \quad \max \quad & \rho\pi^T\hat{v} + (1-\rho)\left(\lambda\hat{v}^+ + (1-\lambda)\hat{v}^-\right) - \rho\pi^T v - (1-\rho)\left(\lambda v^+ + (1-\lambda)v^-\right), \\
\text{s.t.} \quad & \\
& v \geq \hat{v}, \\
& v^- e \leq v \leq v^+ e, \\
& \hat{v}^- e \leq \hat{v} \leq \hat{v}^+ e, \\
& v_{s_i} = v^+, v_{s_j} = \hat{v}^+, v_{s_k} = v^-, v_{s_l} = \hat{v}^-, \\
& v, \hat{v} \in \mathcal{G},
\end{aligned}$$

for $s_i, s_j, s_k, s_l \in S$, $s_i \neq s_k$, $s_j \neq s_l$; and then setting $z = \max_{ijkl}\{z_{ijkl}\}$. Hence, Axiom 4 holds if and only if $z_{ijkl} \leq 0$ for all i, j, k, l. For a given set of indexes i, j, k, l, the dual of the corresponding linear program is

$$
\begin{aligned}
z_{ijkl} = \quad \min \quad & e^T y_2, \\
\text{s.t.} \quad & \\
& -y_1 + y_2 + y_4 - y_6 - \beta_1 e_{s_i} + \beta_3 e_{s_k} = -\rho\pi, \\
& y_1 - y_3 + y_5 - y_7 - \beta_2 e_{s_j} + \beta_4 e_{s_l} = \rho\pi, \\
& -e^T y_4 + \beta_1 = -(1-\rho)\lambda, \\
& -e^T y_5 + \beta_2 = (1-\rho)\lambda, \\
& e^T y_6 - \beta_3 = -(1-\rho)(1-\lambda), \\
& e^T y_7 - \beta_4 = (1-\rho)(1-\lambda), \\
& y_h \in \mathbb{R}^{|S|}, y_h \geq 0, \text{ for all } h \in \{1, \dots, 7\}, \\
& \beta_h \geq 0, \text{ for all } h \in \{1, \dots, 4\}.
\end{aligned}
$$

We claim that the dual problem always has a feasible solution with objective equal to zero, which would imply that $z_{ijkl} \leq 0$ as we wish to show. Hence, set $y_2 = y_3 = 0$, and consider the following cases:

- If $(1-\rho)\lambda \geq 0$ and $(1-\rho)(1-\lambda) \geq 0$, then set $y_5 = y_6 = 0$, $\beta_1 = \beta_4 = 0$, $y_4 = (1-\rho)\lambda e_{s_j}$, $y_7 = (1-\rho)(1-\lambda)e_{s_k}$, $\beta_2 = (1-\rho)\lambda$, $\beta_3 = (1-\rho)(1-\lambda)$, and $y_1 = \rho\pi + (1-\rho)\lambda e_{s_j} + (1-\rho)(1-\lambda)e_{s_k}$.

- If $(1-\rho)\lambda \leq 0$ and $(1-\rho)(1-\lambda) \geq 0$, then set $y_4 = y_6 = 0$, $\beta_2 = \beta_4 = 0$, $y_5 = -(1-\rho)\lambda e_{s_i}$, $y_7 = (1-\rho)(1-\lambda)e_{s_k}$, $\beta_1 = -(1-\rho)\lambda$, $\beta_3 = (1-\rho)(1-\lambda)$, and $y_1 = \rho\pi + (1-\rho)\lambda e_{s_i} + (1-\rho)(1-\lambda)e_{s_k}$.

- If $(1-\rho)\lambda \geq 0$ and $(1-\rho)(1-\lambda) \leq 0$, then set $y_5 = y_7 = 0$, $\beta_1 = \beta_3 = 0$, $y_4 = (1-\rho)\lambda e_{s_j}$, $y_6 = -(1-\rho)(1-\lambda)e_{s_l}$, $\beta_2 = (1-\rho)\lambda$, $\beta_4 = -(1-\rho)(1-\lambda)$, and $y_1 = \rho\pi + (1-\rho)\lambda e_{s_j} + (1-\rho)(1-\lambda)e_{s_l}$.

- If $(1-\rho)\lambda \leq 0$, $(1-\rho)(1-\lambda) \leq 0$, and $s_i \neq s_l$, then set $y_4 = y_7 = 0$, $\beta_2 = \beta_3 = 0$, $y_5 = -(1-\rho)\lambda e_{s_i}$, $y_6 = -(1-\rho)(1-\lambda)e_{s_l}$, $\beta_1 = -(1-\rho)\lambda$, $\beta_4 = -(1-\rho)(1-\lambda)$, and $y_1 = \rho\pi + (1-\rho)\lambda e_{s_i} + (1-\rho)(1-\lambda)e_{s_l}$.

- If $(1-\rho)\lambda \leq 0$, $(1-\rho)(1-\lambda) \leq 0$, $\rho\pi_{s_i} + 1 - \rho \leq 0$, and $s_i = s_l$, then set $y_7 = 0$, $\beta_2 = \beta_3 = 0$, $y_4 = (\rho\pi_{s_i} + (1-\rho)\lambda)e_h$, $y_5 = (\rho\pi_{s_i} + (1-\rho)(1-\lambda))e_{s_i} - (\rho\pi_{s_i} + 1 - \rho)e_{s_{h'}}$, $y_6 = (\rho\pi_{s_i} + (1-\rho)\lambda)e_h - (\rho\pi_{s_i} + 1 - \rho)e_{s_{h'}}$, $\beta_1 = \rho\pi_{s_i}$, $\beta_4 = -(1-\rho)(1-\lambda)$, and $y_1 = \rho\pi + \rho\pi_{s_i}e_{s_i} + (\rho\pi_{s_i} + 1 - \rho)e_{s_{h'}}$, where $s_h \neq s_{h'}$ and $s_h, s_{h'}$ are arbitrary states different from s_i.

- If $(1-\rho)\lambda \leq 0$, $(1-\rho)(1-\lambda) \leq 0$, $\rho\pi_{s_i} + 1 - \rho > 0$, and $s_i = s_l$, then set $y_7 = 0$, $\beta_2 = \beta_3 = 0$, $y_4 = -(1-\rho)(1-\lambda)e_h$, $y_5 = -(1-\rho)\lambda e_{s_i}$, $y_6 = -(1-\rho)(1-\lambda)e_h$, $\beta_1 = -(1-\rho)$, $\beta_4 = -(1-\rho)(1-\lambda)$, and $y_1 = \rho\pi + (1-\rho)e_{s_i}$, where s_h is an arbitrary state different from s_i.

In each of the above cases we exhibit a feasible solution to the dual problem with a zero objective value, so that, as argued before, Axiom 4 holds.

Finally, by noticing that φ is a linear function within any co-extreme class $\mathcal{G}_k$, i.e., $\varphi(\alpha v + (1-\alpha)\hat{v}) = \alpha\varphi(v) + (1-\alpha)\varphi(\hat{v})$ for all $v, \hat{v} \in \mathcal{G}_k$ and $\alpha \in [0,1]$, it is not difficult to see that Axiom 5 also holds. Therefore, Axioms 1-5 are necessary. ■

Additional Representation Results

Proof of Proposition 1: Without loss of generality, we prove this result in state-utility space. Extending the definition of complementary with indifferent extremes (CIE), it is easy to show that two acts $v, \hat{v} \in \mathcal{G}$ are CIE if and only if $v + \hat{v} = (\max(v) + \min(v))e$. Suppose that Axioms 1-6 hold. Then, the relation $\succ^u$ can be represented by a NEU-EU functional φ with parameters π, λ, ρ as in Theorem 2. It follows that for $\rho > 0$, a CIE pair $(v, \hat{v})$ satisfies $v \sim^u \hat{v}$ if and only if $\pi^T v = (\max(v) + \min(v))/2$. Let $(v, \hat{v})$ be a CIE pair satisfying this equation and such that $\max(v) > \min(v)$ (for instance, because $|S| > 3$, there exist $s, s' \in S$ such that $\pi_s + \pi_{s'} \leq 1/2$; so that take $v_s = 1$, $v_{s'} = 0$, and $v_{s''} = x$ for all $s'' \in S \setminus \{s, s'\}$ with $x \in [0,1]$ such that $\pi_s + (1 - \pi_s - \pi_{s'})x = 1/2$). If $\rho = 0$, then any CIE pair $(v, \hat{v})$ trivially satisfies $v \sim^u \hat{v}$, so that again we can choose $\max(v) > \min(v)$. Then, by Axiom 6, $\varphi(\frac{1}{2}v + \frac{1}{2}\hat{v}) \geq \varphi(v)$, which implies

$$(1-\rho)(1/2-\lambda)(\max(v) - \min(v)) \geq 0, \tag{24}$$

and the sufficiency of the axioms follows.

For the converse, assume that $\succ^u$ can be represented by a generalized NEU-EU functional φ with parameters π, λ, ρ satisfying $(1-\rho)(1/2-\lambda) \geq 0$. Theorem 2 implies that Axioms 1-5 hold. For any CIE pair $(v, \hat{v})$ such that $v \sim^u \hat{v}$ inequality (24) is satisfied, which implies $\varphi(\frac{1}{2}v + \frac{1}{2}\hat{v}) \geq \varphi(v)$. Hence, Axiom 6 also holds, and the necessity of the axioms follows. ■

Proof of Proposition 2: Without loss of generality, we prove this result in state-utility space. In this case, two acts $v, \hat{v} \in \mathcal{G}$ are co-maximal if $\overline{S}(v) = \overline{S}(\hat{v})$, and co-minimal if $\underline{S}(v) = \underline{S}(\hat{v})$. Suppose that Axioms 1-5 and 7 hold. Then, the relation $\succ^u$ can be represented by a NEU-EU functional φ with parameters π, λ, ρ as in Theorem 2. Since $|S| > 3$, let s', s'', s''' be three different states in S. Consider the state-utility vectors $v, \hat{v}$ defined as follows:

$$v_s := \begin{cases} 1 & \text{if } s = s', \\ 1/2 & \text{if } s = s'', \\ 0 & \text{if } s = s''', \\ 1/4 & \text{if } s \neq s', s'', s''', \end{cases} \qquad \hat{v}_s := \begin{cases} 1 & \text{if } s = s', \\ 0 & \text{if } s = s'', \\ 1/2 & \text{if } s = s''', \\ 1/4 & \text{if } s \neq s', s'', s'''. \end{cases}$$

Notice that $\overline{S}(v) = \overline{S}(\hat{v}) = \{s'\}$, so that v and $\hat{v}$ are co-maximal. Axiom 7 implies $\varphi\left(\frac{1}{2}v + \frac{1}{2}\hat{v}\right) \geq \frac{1}{2}\varphi(v) + \frac{1}{2}\varphi(\hat{v})$, from which we obtain

$$(1-\rho)(1-\lambda) \geq 0. \tag{25}$$

Now, re-define v and $\hat{v}$ as follows:

$$v_s := \begin{cases} 1/4 & \text{if } s = s', \\ 1/2 & \text{if } s = s'', \\ 1 & \text{if } s = s''', \\ 3/4 & \text{if } s \neq s', s'', s''', \end{cases} \qquad \hat{v}_s := \begin{cases} 1/4 & \text{if } s = s', \\ 1 & \text{if } s = s'', \\ 1/2 & \text{if } s = s''', \\ 3/4 & \text{if } s \neq s', s'', s'''. \end{cases}$$

This time v and $\hat{v}$ are co-minimal. Axiom 7 implies $\varphi\left(\frac{1}{2}v + \frac{1}{2}\hat{v}\right) \leq \frac{1}{2}\varphi(v) + \frac{1}{2}\varphi(\hat{v})$, from which we obtain

$$(1-\rho)\lambda \geq 0. \tag{26}$$

Adding together the left-hand sides of inequalities (25) and (26), it follows that $1-\rho \geq 0$, that is, $\rho \leq 1$. If $\rho = 1$, then we can arbitrarily choose any λ in $[0, 1]$. If $\rho < 1$, then inequalities (25) and (26) imply that $0 \leq \lambda \leq 1$. Therefore, the axioms are sufficient.

For the converse, suppose that $\succ^u$ can be represented by a generalized NEO-

EU functional φ with parameters π, λ, ρ satisfying $\rho, \lambda \in [0,1]$. Theorem 2 implies that Axioms 1-5 hold. If $v \sim^u \hat{v}$ and $v, \hat{v}$ are co-maximal, it follows that $\max\left(\frac{1}{2}v + \frac{1}{2}\hat{v}\right) = \frac{1}{2}\max(v) + \frac{1}{2}\max(\hat{v})$, and $\min\left(\frac{1}{2}v + \frac{1}{2}\hat{v}\right) \geq \frac{1}{2}\min(v) + \frac{1}{2}\min(\hat{v})$. Thus, $\varphi\left(\frac{1}{2}v + \frac{1}{2}\hat{v}\right) \geq \frac{1}{2}\varphi(v) + \frac{1}{2}\varphi(\hat{v}) = \varphi(v)$. Hence, Axiom 7 holds in this case. Similarly, the axiom holds if $v \sim^u \hat{v}$ and $v, \hat{v}$ are co-minimal. Therefore, the axioms are necessary. ■

Proof of VC-Dimension Results

Proposition 7 *There exists a sample of size $|S| - 1$ that can be shattered by $\mathcal{E}$.*

Proof. Let $m := |S| - 1$ and fix a state $\bar{s} \in S$. For each $s \in S \setminus \{\bar{s}\}$, let $v^s := e_s$ and $\hat{v}^s := 0$. Let f_s and $\hat{f}_s$ be such that $M(f_s)\hat{u} = v^s$ and $M(\hat{f}_s)\hat{u} = \hat{v}^s$ for all $s \in S \setminus \{\bar{s}\}$. Let d be the sample m-vector formed by the pairs $(f_s, \hat{f}_s)$ for $s \in S \setminus \{\bar{s}\}$. Let b be any binary m-vector and define $I := \{s \in S \setminus \{\bar{s}\} : b_s = 1\}$ and $J := \{s \in S \setminus \{\bar{s}\} : b_s = 0\}$. Clearly, I and J form a partition of $S \setminus \{\bar{s}\}$ and $|I| + |J| = m$.

If I is nonempty, then define $\pi := \frac{1}{|I|}\sum_{i \in I} e_i$. Notice that $e^T\pi = 1$ and $\pi \geq 0$, so that $\pi \in \Delta(S)$. Moreover, $\pi^T v^s = 1/|I| > 0 = \pi^T \hat{v}^s$ for all $s \in I$, and $\pi^T v^s = 0 = \pi^T \hat{v}^s$ for all $s \in J$. Hence, if $\succ$ is the relation in $\mathcal{E}$ defined by π, we obtain $f_s \succ \hat{f}_s$ for all $s \in I$ and $f_s \not\succ \hat{f}_s$ for all $s \in J$. If I is empty, then define $\pi := e_{\bar{s}}$, so that we have again $\pi \in \Delta(S)$ and $\pi^T v^s = 0 = \pi^T \hat{v}^s$ for all $s \in S \setminus \{\bar{s}\}$, so that $f_s \not\succ \hat{f}_s$ for all $s \in S \setminus \{\bar{s}\}$.

Hence, we can always choose $\pi \in \Delta(S)$ that can separate the sample d according to vector b. Therefore, d can be shattered by $\mathcal{E}$. ■

Proposition 8 *Any sample of size $|S|$ or greater cannot be shattered by $\mathcal{E}$.*

Proof. For arbitrary $m \in \mathbb{N}$, let $d := \left[(f_1, \hat{f}_1), \ldots, (f_m, \hat{f}_m)\right] \in \mathcal{D}_m$, and define $a^i := v^i - \hat{v}^i$, where $v^i = M(f_i)\hat{u}$ and $\hat{v}^i = M(\hat{f}_i)\hat{u}$ for all $i \in \{1, \ldots, m\}$. Let $A := [a^1, \ldots, a^m]$ be the matrix whose columns are the a^i vectors. Let I, J be a partition of $\{1, \ldots, m\}$. Then, there exists a relation $\succ \in \mathcal{E}$ such that $f_i \succ \hat{f}_i$ for all $s \in I$ and $f_j \not\succ \hat{f}_j$ for all $j \in J$ if and only the following linear system is feasible:

$$(*) \qquad A_I^T w \;\geq\; \delta e,$$

$$\begin{aligned} A_J^T w &\leq 0, \\ e^T w &\geq \delta, \\ w &\geq 0, \\ \delta &> 0. \end{aligned}$$

If w is a solution to this system, then it is easy to see that the relation $\succ$ can be represented by $\pi := w/e^T w \in \Delta(S)$. By Farkas' Lemma, system (*) is feasible if and only if the following system is not feasible:

$$(**) \qquad \begin{aligned} -A_I y_I + A_J y_J - \theta e &\geq 0, \\ e_I^T y_I + \theta &= 1, \\ y_I, y_J, \theta &\geq 0. \end{aligned}$$

Now, suppose that $m = |S|$ and $a_1, \ldots, a_m$ are linearly independent vectors. It follows that A is an $m \times m$ nonsingular matrix, so that there exits a vector $q \in \mathbb{R}^m$ such that $Aq = -e$. Let

$$I := \{i : q_i > 0\} \text{ and } J := \{j : q_j \leq 0\}. \tag{27}$$

It follows that $-A_I q_I + A_J(-q_J) - e = 0$. Then, by setting $y_I := \frac{q_I}{e_I^T q_I + 1}$, $y_J := \frac{-q_J}{e_I^T q_I + 1}$, and $\theta := \frac{1}{e_I^T q_I + 1}$, we obtain a solution to system (**), so that system (*) for partition (27) cannot be feasible. Hence, there cannot be a relation $\succ \in \mathcal{E}$ such that $f_i \succ \hat{f}_i$ for all $s \in I$ and $f_j \not\succ \hat{f}_j$ for all $j \in J$ for partition (27).

If $m = |S|$ and $a_1, \ldots, a_m$ are linearly dependent vectors or if $m > |S|$, then there exits a vector $q \in \mathbb{R}^m$ such that $Aq = 0$. It follows that $-A_I q_I + A_J(-q_J) = 0$, where I, J are as in (27). Without loss of generality we can assume that I is nonempty, so that by setting $y_I := \frac{q_I}{e_I^T q_I}$, $y_J := \frac{-q_J}{e_I^T q_I}$, and $\theta := 0$, we obtain another solution to system (**).

Therefore, when $m \geq |S|$ any sample in $\mathcal{D}_m$ cannot be shattered by $\mathcal{E}$. ∎

Corollary 3 *Let $\mathcal{E}' \subset \mathcal{H}$ be such that $\succ \in \mathcal{E}'$ if there exists $w \in \mathbb{R}^S$ such that $e^T w = 1$, and $f \succ g$ if and only if $\sum_{s \in S} w_s U(f(s)) > \sum_{s \in S} w_s U(g(s))$, for all $f, g \in \mathcal{F}$. Then, any sample of size $|S|$ or greater cannot be shattered by $\mathcal{E}'$.*

Proof. The proof is analogous to the proof of Proposition 8, except that we use the linear system:

$$\begin{aligned} A_I^T w &\geq \delta e, \\ A_J^T w &\leq 0, \\ e^T w &\geq \delta, \\ \delta &> 0, \end{aligned}$$

whose alternative system through Farkas' Lemma is:

$$\begin{aligned} -A_I y_I + A_J y_J - \theta e &= 0, \\ e_I^T y_I + \theta &= 1, \\ y_I, y_J, \theta &\geq 0. \end{aligned}$$

Like in Proposition 8, the proof follows from noticing that when the sample vectors $a^1, \ldots, a^m$ are either linearly dependent or linearly independent for $m = |S|$, then the alternative system has a solution and hence, it is not possible to separate the sample vectors by a relation in $\mathcal{E}'$. ■

Proposition 9 *For $|S| > 2$, there exists a sample of size $|S|+1$ that can be shattered by $\mathcal{N}$.*

Proof. Let $m := |S| + 1$, fix a state $s_0 \in S$, and let s_1 be another state not in S. For each $s \in S$, let $v^s := e_s$ and $\hat{v}^s := e/|S|$. Let $v^{s_1} = e - e_{s_0}$ and $\hat{v}^{s_1} := e/2$. Let f_s and $\hat{f}_s$ be such that $M(f_s)\hat{u} = v^s$ and $M(\hat{f}_s)\hat{u} = \hat{v}^s$ for all $s \in S \cup \{s_1\}$. Let d be the sample m-vector formed by the pairs $(f_s, \hat{f}_s)$ for $s \in S \cup \{s_1\}$. Let b be any binary m-vector and define $I := \{s \in S : b_s = 1\}$, and $J := \{s \in S : b_s = 0\}$. Clearly, I and J form a partition of S and $|I| + |J| = |S|$. We consider three cases:

1. I and J are nonempty.

2. I is empty and so, $J = S$.

3. $I = S$ and so, J is empty.

Notice that state s_1 belongs to neither I nor J.

In case #1, define $\pi \in \Delta(S)$ as

$$\pi_s := \begin{cases} \frac{1}{|S|} + \frac{\theta}{|I|} & \text{if } s \in I, \\ \frac{1}{|S|} - \frac{\theta}{|J|} & \text{if } s \in J, \end{cases}$$

for all $s \in S$, where θ is a positive number small enough to ensure that $\pi \geq 0$. Clearly, $e^T\pi = 1$. Also, define $\lambda := 1/|S|$. Then, for $s \in I$, we have $\varphi(v^s) = \frac{1}{|S|} + \rho\frac{\theta}{|I|} > \frac{1}{|S|} = \varphi(\hat{v}^s)$, for any $\rho > 0$. For $s \in J$, we have $\varphi(v^s) = \frac{1}{|S|} - \rho\frac{\theta}{|J|} < \frac{1}{|S|} = \varphi(\hat{v}^s)$, for any $\rho > 0$. If $b_{s_1} = 1$, then we take $\rho = 1$, so that $\varphi(v^{s_1}) = 1 - \frac{1}{|S|} + \theta K > \frac{1}{2} = \varphi(\hat{v}^{s_1})$, for θ small enough, where $K := -1/|I|$ if $s_0 \in I$ and $K := 1/|J|$ if $s_0 \in J$. If $b_{s_1} = 0$, then we take ρ as a positive number close to zero, so that $\varphi(v^{s_1}) = \frac{1}{|S|} + \rho\left(1 - \frac{2}{|S|} + \theta K\right) < \frac{1}{2} = \varphi(\hat{v}^{s_1})$.

In case #2, we set $\pi := e/|S|$ and $\lambda := 1/|S|$. Then, $\varphi(v^s) = 1/|S| = \varphi(\hat{v}^s)$, for all $s \in S$. If $b_{s_1} = 1$, then we take $\rho = 1$, so that $\varphi(v^{s_1}) = 1 - 1/|S| > 1/2 = \varphi(\hat{v}^{s_1})$, and if $b_{s_1} = 0$, then we take $\rho = 0$, so that $\varphi(v^{s_1}) = 1/|S| < 1/2 = \varphi(\hat{v}^{s_1})$.

In case #3, we set $\pi := e/|S|$ and $\rho := 0$. Then, $\varphi(v^s) = \lambda > 1/|S| = \varphi(\hat{v}^s)$, for all $\lambda > 1/|S|$ and $s \in S$. If $b_{s_1} = 1$, then we take $\lambda = 1$, so that $\varphi(v^{s_1}) = 1 > 1/2 = \varphi(\hat{v}^{s_1})$, and if $b_{s_1} = 0$, then we take $\lambda = 1/2$, so that $\varphi(v^{s_1}) = 1/2 = \varphi(\hat{v}^{s_1})$.

Hence, we can always choose $(\pi, \lambda, \rho) \in \Delta(S) \times [0, 1] \times [0, 1]$ that can separate the sample d according to vector b. Therefore, d can be shattered by $\mathcal{N}$. ■

Proposition 10 *Any sample of size $|S| + 2$ or greater cannot be shattered by $\mathcal{N}_1$.*

Proof. Let $\succ$ be a relation in $\mathcal{N}_1$. Then, there exist $\pi \in \mathbb{R}^S$, $e^T\pi = 1$, and $\rho, \lambda \in \mathbb{R}$ such that for any two acts $f, \hat{f} \in \mathcal{F}$ we have $f \succ \hat{f}$ if and only if

$$\rho\pi^T v + (1-\rho)\left(\lambda \max(v) + (1-\lambda)\min(v)\right) > \\ \rho\pi^T \hat{v} + (1-\rho)\left(\lambda \max(\hat{v}) + (1-\lambda)\min(\hat{v})\right), \quad (28)$$

where $v = M(f)\hat{u}$ and $\hat{v} = M(\hat{f})\hat{u}$. Now, let $S' := S \cup \{s_0, s_1\}$, where s_0 and s_1 are different states not in S. Define $w \in \Delta(S')$ as

$$w_s := \begin{cases} \rho\pi_s & \text{if } s \in S, \\ (1-\rho)\lambda & \text{if } s = s_0, \\ (1-\rho)(1-\lambda) & \text{if } s = s_1, \end{cases}$$

for all $s \in S'$. Then, inequality (28) becomes $w^T v' > w^T \hat{v}'$, where $v' := [v, \max(v), \min(v)]$ and $\hat{v}' := [\hat{v}, \max(\hat{v}), \min(\hat{v})]$. If f' and $\hat{f}'$ are augmented acts defined in S' such that $M(f')\hat{u} = v'$ and $M(\hat{f}')\hat{u} = \hat{v}'$, then $f \succ \hat{f}$ if and only if $f' \succ' \hat{f}'$, where $\succ'$ is the relation in $\mathcal{E}'$ from Corollary 3 defined by w.

If $m = |S|$ and $a_1, \ldots, a_m$ are linearly dependent vectors or if $m > |S|$, then there exists a nonzero vector $q \in \mathbb{R}^m$ such that $Aq = 0$. It follows that $-A_I q_I + A_J(-q_J) = 0$, where I, J are as in (27). Without loss of generality we can assume that I is nonempty, so that by setting $y_I := \frac{q_I}{e_I^T q_I}$, $y_J := \frac{-q_J}{e_I^T q_I}$, and $\theta := 0$, we obtain another solution to system (**).

Hence, if an arbitrary sample $d := \Big[(f_1, \hat{f}_1), \ldots, (f_m, \hat{f}_m)\Big] \in \mathcal{D}_m$ can be shattered by $\mathcal{N}_1$, then the corresponding sample of augmented acts $d' := \Big[(f_1', \hat{f}_1'), \ldots, (f_m', \hat{f}_m')\Big]$ can also be shattered by $\mathcal{E}'$ on the augmented state space S'. However, because $|S'| = |S| + 2$, Corollary 3 implies that d' cannot be shattered by $\mathcal{E}'$ for $m \geq |S| + 2$. Therefore, d cannot be shattered by $\mathcal{N}_1$ for $m \geq |S| + 2$ and the result follows. ■

Proof of Theorem 3: Proposition 7 implies that $\mathrm{VC}(\mathcal{E}) \geq |S| - 1$, whereas Proposition 8 implies that $\mathrm{VC}(\mathcal{E}) < |S|$. Therefore, it follows that $\mathrm{VC}(\mathcal{E}) = |S| - 1$.

For $|S| > 2$, Proposition 9 implies that $\mathrm{VC}(\mathcal{N}) \geq |S| + 1$, whereas Proposition 10 implies that $\mathrm{VC}(\mathcal{N}_1) < |S| + 2$. Hence, we obtain $|S| + 1 \leq \mathrm{VC}(\mathcal{N}) \leq \mathrm{VC}(\mathcal{N}_1) < |S| + 2$ Therefore, it follows that $\mathrm{VC}(\mathcal{N}) = \mathrm{VC}(\mathcal{N}_1) = |S| + 1$.

Finally, to show that $\mathrm{VC}(\mathcal{N}_0) = 1$, we first show that there is a sample vector in $\mathcal{D}_1$ that can be shattered. Let $\bar{s}$ be a fixed state in S. Let $v := e_{\bar{s}}$, $\hat{v} := \frac{1}{|S|} e$, and $b \in \{0, 1\}$. Let f and $\hat{f}$ be such that $M(f)\hat{u} = v$ and $M(\hat{f})\hat{u} = \hat{v}$, and $d := (f, \hat{f})$. If $b = 1$, then let $\succ$ be the Hurwicz preference corresponding to $\lambda = 1$. Then, we have $\nu(f) = 1 > \frac{1}{|S|} = \nu(\hat{f})$. If $b = 0$, then let $\succ$ be the Hurwicz preference corresponding to $\lambda = 0$. Then, we have $\nu(f) = 0 < \frac{1}{|S|} = \nu(\hat{f})$. Hence, d can be shattered by $\mathcal{N}_0$.

Next, we show that there cannot be a sample in $\mathcal{D}_2$ shattered by $\mathcal{N}_0$. Let $d := [(f_1, \hat{f}_1), (f_2, \hat{f}_2)] \in \mathcal{D}_2$. By letting parameter λ vary in the interval $[0, 1]$, it follows that the mappings $g_1(\lambda) := \nu(f_1; \lambda)$, $\hat{g}_1(\lambda) := \nu(\hat{f}_1; \lambda)$, $g_2(\lambda) := \nu(f_2; \lambda)$, and $\hat{g}_2(\lambda) := \nu(\hat{f}_2; \lambda)$ are linear functions on $[0, 1]$. Notice that to shatter d it is necessary and sufficient to partition interval $[0, 1]$ into at least four nonempty regions such that each region corresponds to each of the four possible conditions $g_1(\lambda) - \hat{g}_1(\lambda) >$ or ≤ 0 and $g_2(\lambda) - \hat{g}_2(\lambda) >$ or ≤ 0. If such a partition exists, then d can be shattered by taking the Hurwicz preferences corresponding to four values of λ taken respectively from each region in the partition. However, because $g_1, \hat{g}_1, g_2, \hat{g}_2$ correspond to straight lines, there is at most one point λ_1 such that $g_1 - \hat{g}_1$ changes sign in $[0,1]$, and there is at most one point λ_2 such that $g_2 - \hat{g}_2$ changes sign in $[0,1]$. Taken together, the points λ_1 and λ_2 define at most three nonempty regions in $[0,1]$, where $g_1(\lambda) - \hat{g}_1(\lambda) >$ or ≤ 0 and $g_2(\lambda) - \hat{g}_2(\lambda) >$ or ≤ 0. Therefore, it is not possible to shatter d in this case, and we must have $\mathrm{VC}(\mathcal{N}_0) = 1$. ■

Proof of Proposition 3: The dual problem of (P) is as follows:

$$
\begin{array}{lll}
\text{(D)} & \min & \theta, \\
& \text{s.t.} & \\
& & -A_I y_I + A_J y_J + \theta e \geq 0, \\
& & -g_I^T y_I + g_J^T y_J + \theta \geq 0, \\
& & -h_I^T y_I + h_J^T y_J + \theta \geq 0, \\
& & e_I^T y_I \geq 1, \\
& & y_I, y_J \geq 0.
\end{array}
$$

Since we can choose θ to be an arbitrarily large positive number, problem (D) is always feasible. Hence, if problem (P) is feasible, its objective function will be bounded from above and (P) will have an optimal solution.

If (P) has a solution $(\alpha, \beta, \sigma, w)$ with $\sigma > 0$, then let $(\pi, \lambda, \rho) \in \Delta(S) \times [0,1] \times [0,1]$ be defined using (α, β, w) as indicated in the statement of this proposition, and $\succ$ the relation in $\mathcal{N}$ defined by (π, λ, ρ). It follows that for each $i \in I$ we have $\rho(v^i - \hat{v}^i)^T \pi + (1-\rho)\,[\lambda(\max(v^i) - \max(\hat{v}^i)) + (1-\lambda)(\min(v^i) - \min(\hat{v}^i))] > 0$, which implies that $f_i \succ \hat{f}_i$. Similarly, for each $i \in J$,we have $\rho(v^i - \hat{v}^i)^T \pi + (1-\rho)[\lambda(\max(v^i) - \max(\hat{v}^i)) + (1-\lambda)(\min(v^i) - \min(\hat{v}^i))] \leq 0$, which implies that $f_i \not\succ \hat{f}_i$. Therefore, under this relation we have $f_i \succ \hat{f}_i \Leftrightarrow b_i = 1$ for $i = 1, \ldots m$.

For the converse, suppose that $f_i \succ \hat{f}_i \Leftrightarrow b_i = 1$ for $i = 1, \ldots m$ for some $\succ$ in $\mathcal{N}$. Let $(\pi, \lambda, \rho) \in \Delta(S) \times [0,1] \times [0,1]$ be a corresponding parameter vector defining $\succ$. Then, it is easy to see that if we define $w := \rho\pi$, $\alpha := (1-\rho)\lambda$, and $\beta := (1-\rho)(1-\lambda)$, then we can obtain a feasible solution to (P) with $\sigma > 0$, and the result follows. ■