MISSPECIFICATION AVERSE PREFERENCES*

Alfonso Maselli†

August 20, 2024

INCOMPLETE AND VERY PRELIMINARY DRAFT
PLEASE, DO NOT CIRCULATE

Abstract. We study a decision maker who approaches an uncertain decision problem by formulating a set of plausible probabilistic models of the environment but is aware that these models are only stylized and incomplete approximations. We introduce the concept of a best-fit map that identifies the most suitable model within this potentially misspecified set based on observable data. Building on this, we develop an axiomatic foundation for preferences that are averse to misspecification. In particular, we introduce a novel criterion that discriminates between aversion to misspecification and attitudes toward model ambiguity. First, conditional on a model having the best fit, the decision maker forms a misspecification-robust evaluation by considering a range of models in proximity to the best-fit one. Then, she aggregates these robust evaluations via a monotone and quasiconcave aggregator incorporating uncertainty about what model is the best approximation of the environment.

1. Introduction

Economic agents often employ simplified and stylized descriptions of the complex environment they face in order to help guide their decisions. This implies that model

*I am indebted to David Dillenberger, George Mailath, and Kevin He for their valuable guidance. I am also grateful to Aislinn Bohren, Simone Cerreia-Vioglio, Roberto Corrao, Giacomo Lanzani, Fabio Maccheroni, Julien Manili, Massimo Marinacci, Andrew Postlewaite, Sofia Teles, and Juuso Toikka for helpful comments and suggestions.
†Department of Economics, University of Pennsylvania. Email: masellia@sas.upenn.edu.

misspecification is a pervasive phenomenon affecting many decision problems. For example, a policymaker might have an incorrect description of how the economy would respond to a fiscal or monetary stimulus, or a company's marketing department might have a wrong assessment of how demand would react to changes in the price of a product. As a result, a growing literature studies the implications of using misspecified models in the context of decision making and strategic interaction. Starting with Esponda and Pouzo (2016), many papers have examined the asymptotic behavior of actions and beliefs when agents take repeated decisions in a stochastic environment of which they have a possibly incorrect or only partial understanding (see, for instance, Frick, Iijima, and Ishii, 2022; Fudenberg, Lanzani, and Strack, 2021). A related strand of the literature has focused on studying whether misspecification is asymptotically persistent. For example, Ba (2021), Fudenberg and Lanzani (2023), and He and Libgober (2021) provide conditions to identify whether and which misspecifications will persist in the long run. These papers suggest that misspecification matters in shaping agents' behavior and beliefs and that it is a persistent phenomenon, even when agents collect many observations generated by the true data-generating process. A common assumption in this literature is that once agents have settled on using a specific statistical model of the environment, they disregard the possibility of it being misspecified and act in a fully Bayesian fashion, as they evaluate alternative actions by computing their expected utility with respect to their model. However, sophisticated enough agents should realize that their model is only a simplified approximation of reality. As suggested by Hansen and Sargent (2001), a decision maker who is concerned with acting on the basis of an incorrectly formulated model should make decisions that are *robust*; that is, policies that do not depend on the fine details of their reference model, but work reasonably well across all models that are perturbations of that reference model. Following this idea, first axiomatic treatments of decision criteria featuring misspecification aversion have been proposed by Cerreia-Vioglio et al. (2020) and Lanzani (2022).

In this paper, we provide an axiomatic foundation of a general class of preferences that are averse to the possibility of misspecification. We introduce a way of meaningfully disentangling misspecification aversion from the more commonly studied aversion to model ambiguity and show that these are captured independently by two different elements of decision criterion obtained in the main representation result. In particular, we study a decision maker who faces a decision problem in a generalized version of the Anscombe-Aumann setting. The uncertainty is captured by a set of states of the world Ω and the decision maker needs to choose an act f that maps states of the world to outcomes. The decision maker does not know the true stochastic process governing the environment, but she has statistical information in her possession. This is given by a set $\mathcal{M}$ of distributions over states of the world. Following Cerreia-Vioglio et al. (2020) and Hansen and Sargent (2022), we interpret distributions in $\mathcal{M}$ as being explicitly motivated on the basis of scientific knowledge or empirical considerations. This is in keeping with the classical setup of Wald (1950), according to which the set $\mathcal{M}$ would be interpreted as a collection of alternative hypotheses regarding the data-generating process (DGP) under study. Unlike Wald (1950), we allow for the possibility that the set $\mathcal{M}$ does not include the true DGP. A decision maker who is aware that models are only imperfect and stylized descriptions of the real environment might become concerned that, in fact, no structured model is an accurate approximation of the DGP. In order to differentiate between the ambiguity about which structured model is the best approximation to the DGP and the decision maker's perception of and concern about misspecification, we follow the approach in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2013) and formalize the notion of the missing information the decision maker would need to determine the best structured model via the idea of sufficient statistics and information (Dynkin, 1978). In particular, we assume that there exists a *best-fit* map $\mathfrak{q}: \Omega \to \mathcal{M}$ measurable with respect to a sigma-algebra $\mathcal{A}$ that encodes the sufficient information to determine the best structured model. This can be seen as the decision maker having a possibly misspecified statistical procedure that would allow the decision maker to identify the best approximation in $\mathcal{M}$ to the true DGP given different realizations of the state of the world. After observing the additional information in $\mathcal{A}$, she would be able to dissipate the model ambiguity within this misspecified procedure, but she would not be able to infer whether the statistical procedure itself is misspecified or not. In other words, $\mathfrak{q}$ can be interpreted as an estimator of the DGP that, because of the possibility of using a misspecified set of models, converges almost surely not necessarily to the true probability law, but to its closest approximation (see Berk, 1966) among the structured models.

Endowed with this structure, we provide an axiomatic foundation of a general misspecification averse decision criterion. First of all, we characterize preferences conditional on a structured model $m \in \mathcal{M}$ being the best-fit model and show that they are represented by the misspecification-robust criterion

$$V^m(f) = \min_{p \in \Delta(\Omega)} \{\mathbb{E}_p[u(f)] + c(p,m)\} \tag{1}$$

where u is a utility over outcomes and $c(\cdot, m)$ is an index of misspecification aversion. That is, even conditional on observing sufficient information to determine that m is the best structured model, the decision maker would not completely trust it out of misspecification concerns. Therefore, in evaluating an act f, she would also take into account other distributions p outside of $\mathcal{M}$ that are not too far apart from m. The index $c(\cdot, m)$ captures exactly the decision maker's confidence in the structured model m. An important special case is given by $c(\cdot, m) = \lambda R(\cdot||m)$, where R is the relative entropy and $\lambda > 0$ is a parameter of misspecification aversion. When the decision maker's concern for misspecification is high, λ is low, and therefore, she would give preference to acts that perform robustly well across a larger set of models around m. The idea that even if aware of the possibility of misspecification, the decision maker still puts substantive trust in the set of structured models is captured by an axiom of consistency. If the misspecification-robust evaluation of an act f dominates that of another act g unanimously according to all structured models, the decision maker should prefer f to g. This axiom, which is the misspecification averse analogue of the consistency requirement in Cerreia-Vioglio et al. (2013), allows us to show that the preferences of a misspecification averse decision maker are represented by aggregating together the misspecification-robust evaluations:

$$V(f) = \hat{I}\left((V^m(f))_{m\in\mathcal{M}}\right) = \hat{I}\left(\min_{p\in\Delta(\Omega)}\{\mathbb{E}_p[u(f)] + c(p,\cdot)\}\right) \tag{2}$$

where $\hat{I} : \mathbb{R}^{\mathcal{M}} \to \mathbb{R}$ is a monotone and quasiconcave aggregator capturing the decision maker's attitudes towards the ambiguity regarding what structured model is the best-fit one. After the main representation result, we perform a comparative statics exercise that clarifies how the decision criterion above disentangles misspecification aversion from aversion to model ambiguity. In particular, we show that we can rank two decision makers in terms of their degree of misspecification aversion by only comparing their misspecification index c (without imposing any mutual restrictions on their aggregators $\hat{I}$). In particular, decision maker 1 is more misspecification averse than decision maker 2 if and only if the index of misspecification aversion of the first one is always lower than that of the second one for each structured model; that is, $c_1(\cdot, m) \leq c_2(\cdot m)$ for all $m \in \mathcal{M}$. Similarly, we show that we can rank decision makers in terms of their attitudes toward model ambiguity by only comparing their aggregator $\hat{I}$ (without imposing any mutual restrictions on their misspecification aversion indexes). In particular, decision maker 1 is more averse to model ambiguity than decision maker 2 if and only if $\hat{I}_1 \leq \hat{I}_2$. Since we can interpret the aggregator as a certainty equivalent, this amounts to saying the first decision maker is more model ambiguity averse if she is willing to accept lower certainty equivalents than the second to eliminate the ambiguity regarding the identity of the best-fit model.

Different assumptions regarding the linearity properties of the preferences will characterize specific shapes of the aggregator $\hat{I}$. In particular, we provide an axiomatization of two important cases. First, we show that when the decision maker confronts the uncertainty regarding the identity of the best-fit model according to the subjective expected utility tenets, she then aggregates the misspecification-robust evaluations in a Bayesian fashion. In this case, the decision maker forms a subjective prior μ over the set of structured models $\mathcal{M}$ and takes a quasi-mean of the robust evaluations using this subjective belief and an index ϕ capturing her attitudes towards the uncertainty about the best-fit model:

$$V_{\phi,\mu}(f) = \int_{\mathcal{M}} \phi\left(\min_{p\in\Delta(\Omega)} \{\mathbb{E}_p[u(f)] + c(p,m)\}\right) d\mu(m). \tag{3}$$

In particular, if the decision maker is neutral towards model ambiguity and shows a uniform concern for misspecification, this criterion becomes the average robust control representation axiomatized by Lanzani (2022).[1] It is also worth noticing how this criterion collapses to the well-known smooth ambiguity model of Klibanoff et al. (2005) under the assumption of misspecification neutrality; that is, when $c(\cdot, m)$ assigns an infinite penalization to any probability model different from m itself. Moreover, we show that if the decision maker is cautious and evaluates the uncertainty about the best-fit model according to a worst-case scenario approach, then the aggregator takes on a maxmin form and we obtain the criterion proposed by Cerreia-Vioglio et al.

[1]To be precise, we would also need to impose that the conditional misspecification-robust evaluations are the multiplier preferences proposed by Hansen and Sargent (2001) and axiomatized by Strzalecki (2011).

(2020):

$$V_{min}(f) = \min_{p\in\Delta(\Omega)} \left\{ \mathbb{E}_p[u(f)] + \min_{m\in\mathcal{M}} c(p,m) \right\}. \tag{4}$$

Discussion and Literature Review. In this paper we revisit the framework introduced in Cerreia-Vioglio et al. (2013) and allow to incorporate misspecification aversion in the preferences of a decision maker who uses exogenous, statistical information to inform her choices. The Dynkin space structure $(\Omega, \mathcal{G}, \mathcal{P})$ in Cerreia-Vioglio et al. (2013) can be interpreted as the decision maker being sure she has available a statistical framework that is correctly specified. The sufficient σ-algebra $\mathcal{A}$ represents for the decision maker exactly the missing information she would need to identify the true probabilistic model of the world. Therefore, upon observing this information, the decision maker should just evaluate acts according to their expected utility with respect to $P \in \mathcal{P}$. Therefore, the set of models induces an objectively rational preference; that is a dominance relation in terms of expected utility certainty equivalents with respect to models in $\mathcal{P}$:

$$\forall P \in \mathcal{P},\ \int_S f dP \succsim \int_S g dP.$$

In our case, however, even after observing the missing information sufficient to pin down a unique best-fit model $m \in \mathcal{M}$, the decision maker, out of misspecification concerns, would only trust m to be the best approximation to the DGP among the structured models, but not necessarily the correct distributions of states of the world. Therefore, our objectively rational dominance relation is given by the unanimity criterion with respect to the preferences conditional on each structured model $m \in \mathcal{M}$. Since the possibility of misspecification implies that uncertainty about the true probabilistic model is not resolved even after observing the additional information (only the ambiguity regarding the identity of the best approximation is), such conditional preferences need not be expected utility with respect to the structured model but can still display a preference for robustness across models that are in a vicinity of m. From a mathematical perspective, Cerreia-Vioglio et al. (2013) show that consistency of the subjective preferences with the objectively rational dominance relation imply that any utility representation only depends on the profile of expected utility evaluations $(\mathbb{E}_P[u(f)])_{P\in\mathcal{P}}$. This allows their preferences to be represented via an aggregator of the map $P \mapsto \mathbb{E}_P[u(f)]$. In our case, due to misspecification aversion, in Theorem 4 we show that the representation of our class of misspecification averse preferences only depends on the profile of misspecification robust conditional evaluations $(\min_p \mathbb{E}_p[u(f)] + c(p,m))_{m\in\mathcal{M}}$, so that the representation can be expressed as a certainty equivalent $\hat{I}$ of the map $m \mapsto \min_p \mathbb{E}_p u(f) + c(p,m)$. The fact that this map is no longer linear in the models $m \in \mathcal{M}$ is the main technical difficulty that we deal with in this paper.[2] In particular, we show that also in our case properties of preferences over acts can be translated into properties of the certainty equivalent $\hat{I}$ without having to resort to second-order acts.

This paper is closely related to the literature on decision criteria that incorporate misspecification aversion. There are a few papers proposing axiomatizations of such preferences. Cerreia-Vioglio et al. (2020) axiomatize the criterion (4) in a two-preference setup. The decision maker has both a mental preference, assumed to be an incomplete variational preference, and a behavioral preference that is, instead, complete but satisfies independence only on constant acts. These two preferences are connected to each other via two axioms that originated in the seminal work of Gilboa et al. (2010). The first is a consistency requirement that the behavioral preferences always agree with the mental ones. The second is that the decision maker exercises caution; that is, if the mental preference is not confident enough to rank an uncertain

[2] In this respect, this paper is also related to Mu et al. (2021). In a different context, they show that monotone additive statistics can be represented as averages of CARA certainty equivalents.

act over a deterministic one, then the behavioral preferences rank the deterministic outcome over the uncertain one (in other words, the decision maker picks the safe alternative whenever in doubt according to her mental preferences). Moreover, the two preferences are informed by the set of structured models via a coherence requirement analogous to the one given in this paper. Lanzani (2022) also adopts the view of Cerreia-Vioglio et al. (2013) by considering states of the world that describe both the realization of the payoff relevant state and the distribution over such payoff states. They assume that the decision maker has variational preferences and obtain the average robust criterion (3) by imposing that preferences on bets over models satisfy the sure thing principle and uncertainty neutrality (thus obtaining an affine ϕ). Moreover, they propose axioms that characterize the asymptotic behavior of the index of misspecification concern when the decision maker's preferences evolve in reaction to the arrival of new information. We show (Theorems 6 and 7) that the criteria introduced by Lanzani (2022) and Cerreia-Vioglio et al. (2020) both fall within the general class of misspecification averse preferences studied in this paper and represent two opposite ends of the spectrum; the average robust criterion is neutral towards model ambiguity, while the maxmin criterion displays an extreme form of model ambiguity aversion. One contribution of our paper is to allow more flexible attitudes toward model ambiguity while proposing a way to disentangle those from the degree of misspecification aversion. This is reflected in the fact that the representation parameters capturing model misspecification aversion (the index c) and model ambiguity aversion (the aggregator $\hat{I}$) are independent of each other.

This paper is also related to the recent axiomatization by Denti and Pomatto (2022) of identifiable smooth ambiguity preferences. In a purely subjective framework, they find conditions under which the preferences are represented by the smooth ambiguity criterion, where the beliefs involved in the representation are identifiable; that is, they are completely orthogonal for some kernel κ. In this paper, we take the reverse approach. We start with a decision maker employing an exogenous statistical model and then impose conditions so that the subjective preferences are informed by this statistical model. This approach might be more natural in discussing the issue of misspecification, which is usually defined in the context of parametric statistical models which might fail to include the true parameter. However, it would be an interesting exercise to extend the techniques of Denti and Pomatto (2022) to identify in a purely subjective framework the perceived misspecification and misspecification concern revealed by the preferences.

The rest of the paper is structured as follows. Section 2 lays out the decision framework and the notions of structured space and best-fit map. Section 3 introduces and discusses the axioms characterizing the misspecification averse preferences. Section 4 states and discusses the representation results. Section 5 concludes. All proofs can be found in the Appendix.

2. Decision Framework

We begin by describing the decision environment faced by the decision maker (DM). Uncertainty is described by a state space Ω endowed with a countably generated σ-algebra $\mathcal{G}$. Fix X to be the space of *consequences*, a non-empty, convex subset of a linear space. The decision maker needs to choose *simple acts*, that is simple functions $f : \Omega \to X$ mapping states to consequences that are measurable with respect to $\mathcal{G}$. Denote by $\mathcal{F}$ the set of all such simple acts. As usual, we can embed X in $\mathcal{F}$ by abusing notation and denoting with $x \in X$ the constant act yielding consequence x in each state of the world $\omega \in \Omega$. We can define the operation of convex combination in $\mathcal{F}$ in the natural way: for all $f, f' \in \mathcal{F}$ and for all $\alpha \in [0, 1]$,

$$(\alpha f + (1-\alpha) f')(\omega) \coloneqq \alpha f(\omega) + (1-\alpha) f'(\omega)$$

for all $\omega \in \Omega$. The affine structure of X implies that $\alpha f + (1-\alpha) f' \in \mathcal{F}$. Then, the set of simple acts $\mathcal{F}$ with the operation just defined is a mixture space, as defined by Herstein and Milnor (1953). We introduce a few useful pieces of notation. Given any $E \in \mathcal{G}$ and simple acts $f, g \in \mathcal{F}$, denote by fEg the act taking value $f(\omega)$ if $\omega \in E$ and value $g(\omega)$ if $\omega \in \Omega \setminus E$. Moreover, if $\mathcal{E}$ is a sub-σ-algebra of $\mathcal{G}$, denote by $\mathcal{F}(\mathcal{E})$ the subset of simple acts in $\mathcal{F}$ that are measurable with respect to $\mathcal{E}$.

We model the decision maker's subjective preferences on the set of simple acts $\mathcal{F}$ via a binary relation $\succsim$ and we denote by $\succ$ and $\sim$ respectively the asymmetric and symmetric part of $\succsim$. We say that an even $E \in \mathcal{G}$ is *null* if for all acts $f, f' \in \mathcal{F}$, $f|_{\Omega \setminus E} = \tilde{f}|_{\Omega \setminus E}$ implies that $f \sim f'$. We say that an event is *nonnull* if it is not null.

2.1. **Structured Models and Best-Fit Map.** We denote by $\Delta := \Delta(\Omega, \mathcal{G})$ and $\Delta^\sigma := \Delta^\sigma(\Omega, \mathcal{G})$ respectively the space of finitely and countably additive probability measures on $(\Omega, \mathcal{G})$. Moreover, we endow Δ^σ with the natural σ-algebra $\mathcal{D}$ generated by the family of evaluations maps and any subset of Δ^σ, with its relative σ-algebra.[3]

We assume that the decision maker has constructed a set $\mathcal{M} \subseteq \Delta^\sigma$ of probability distributions over states of the world that, due to external information and considerations, she believes are plausible descriptions of the uncertain environment she is facing. Following the terminology established in Hansen and Sargent (2022) and Cerreia-Vioglio et al. (2020), we call the probabilities distributions in this set *structured models*. We interpret structured models as being statistical descriptions of the environment that are based on substantive motivations, like scientific theories and evidence. Following Box and Cox's idea that models are only approximations, we do not assume that the set of structured models includes the data-generating process (DGP), that is, the true probability law governing state uncertainty. Moreover, we allow for the possibility that the decision maker is aware of this fact; that is, the

[3]Appendix A provides rigorous definitions of the mathematical concepts and details regarding the notation.

decision maker perceives the possibility that her set of structured models might be *misspecified*.

In this paper, we want to discern between ambiguity about which structured model is the best approximation to the DGP and concern about misspecification; that is, the fact that no structured model is an accurate approximation of the DGP. Uncertainty about models is usually motivated in terms of "lack of information" preventing the the DM from selecting the best one. Following Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2013),[4] we formalize this missing information via the idea of sufficient statistics and information (Dynkin, 1978).

Definition 1: We say that a measurable space and the posited set of structured models $(\Omega, \mathcal{G}, \mathcal{M})$ form a *structured space* if $\mathcal{M}$ is measurable and we can find a *best-fit map* $\mathfrak{q} : \Omega \to \Delta^\sigma$ and a *sufficient* sub σ-algebra $\mathcal{A}$ of $\mathcal{G}$, satisfying the following properties:

(i) $\mathcal{A}$ is the σ-algebra generated by $\mathfrak{q}$,

(ii) $m(\{\omega \in \Omega : \mathfrak{q}(\omega) = m\}) = 1$ for all $m \in \mathcal{M}$.

Note that in our definition of structured space, we allow the best-fit map to select also probability models that are not in $\mathcal{M}$. However, if we denote by Ω_0 the set of states of the world $\omega \in \Omega$ such that $\mathfrak{q}(\omega) \in \mathcal{M}$, we can see that condition (*ii*) implies that $m(\Omega_0) = 1$ for all structured models $m \in \mathcal{M}$. The requirement that each structured model $m \in \mathcal{M}$ is selected by the best-fit map with probability one[5] is equivalent to the notion of sufficient statistics introduced by Dynkin (1978) and is related to the strong law of large numbers. We interpret this framework as follows. Suppose that the well-specified description of the environment is given by the set of models $\mathcal{P}$ and a map $p : \omega \mapsto p^\omega \in \Delta^\sigma$ such that $P(\{\omega \in \Omega : p^\omega = P\}) = 1$ for

[4]See also Amarante (2009), Al-Najjar and De Castro (2014), Epstein and Seo (2010), and Klibanoff et al. (2014) for related approaches.

[5]In mathematics and probability, this property is what is known as complete orthogonality of the set of probability models $\mathcal{M}$. See, for example, Mauldin et al. (1983) and Weis (1984).

all $P \in \mathcal{P}$. The interpretation is that the realization of the state of the world also pins down what is the true DGP p^ω, so that sigma-algebra of events $\mathcal{A}$ that makes p^ω captures the sufficient information to determine what is the true probability law over states. Moreover, the statement that $p^\omega = P$ with probability one according to P is a requirement that the correct description of the environment is not contradictory; that is, whenever P is the true DGP, then it is selected with probability one by the map p^ω.[6] However, the decision maker posits a misspecified set of models $\mathcal{P}_0$ that does not necessarily include all models in $\mathcal{P}$[7]. Now, suppose the decision maker observed the missing information in $\mathcal{A}$ that would be sufficient to infer $P \in \mathcal{P}$. However, since the decision maker has posited a misspecified set of models, the insight from the classical result of Berk (1966) suggests she would select from $\mathcal{P}_0$ the closest model to P; that is, she would select the unique minimizer $q^*(P) \in \mathcal{P}_0$ solving $\min_{q \in \mathcal{P}_0} R(q||P)$, where $R(\cdot||\cdot)$ is the relative entropy.[8] Then, if we define $\mathfrak{q}(\omega) = q^*(p^\omega)$[9] and $\mathcal{M} = \{m \in \Delta^\sigma : \exists P \in \mathcal{P},\ m = q^*(P)\}$, we can notice that for all $m \in \mathcal{M}$ we would, indeed, have that $m\left(\{\omega : \mathfrak{q}(\omega) = m\}\right) = 1$[10]. It is in this sense that we interpret $\mathfrak{q}$ as a *best-fit map* and the information in $\mathcal{A}$ as the sufficient information to determine the best approximation of the DGP among those in $\mathcal{M}$. That is, if the decision maker were able to observe ω, she would infer that the model $m_\omega = \mathfrak{q}(\omega)$ is the model that closest resembles the true DGP.

[6] We can see the analogy to the strong law of large numbers if we interpret each ω as the realization of an infinite sequence of random variables and p^ω as the limit of a consistent estimator.

[7] Assume that $\mathcal{P}_0$ is compact and convex and that for each $P \in \mathcal{P}$, there exists a model $q \in \mathcal{P}_0$ such that q is absolutely continuous with respect to P (written $q \ll P$).

[8] Recall that for every $q, p \in \Delta^\sigma$, $R(q||p) = \int_\Omega \ln \frac{dq}{dp} dq$ if $q \ll p$ and equal to ∞ otherwise. Notice that a minimizer exists since we are assuming that $\mathcal{P}_0$ is compact and contains at least one $q \ll P$ and it is unique since R is strictly convex in its first argument and $\mathcal{P}_0$ is assumed to be a convex set.

[9] Notice that by the measurable maximum theorem, $p \mapsto \arg\max_{q \in \mathcal{P}_0 : q \ll p} R(q||p)$ is a measurable function, so that $\mathfrak{q}$ so defined is $\mathcal{A}$-measurable.

[10] For each $m \in \mathcal{M}$, there exists $P \in \mathcal{S}$ such that $m = q^*(P)$, so that $\{\omega : p^\omega = P\} \subseteq \{\omega : \mathfrak{q}(\omega) = q^*(p^\omega)\}$ and, therefore, $P(\Omega \setminus \{\omega : \mathfrak{q}(\omega) = m\}) \leq P(\Omega \setminus \{\omega : p^\omega = P\}) = 0$. Since $m \ll P$, it then must be the case that $m(\Omega \setminus \{\omega : \mathfrak{q}(\omega) = m\}) = 0$.

Example 1 (Exchangeability): Suppose S is an underlying finite set of *contemporaneous states* and assume that at each time period $t \in \mathbb{N}$, the uncertainty is described by the realization of a contemporaneous state $s_t \in S$. Then, a state of the world is an infinite sequence of realizations from S, and the state space is given by the sequence space $\Omega = S^{\mathbb{N}}$. In this case, the relevant σ-algebra $\mathcal{G}$ is that generated by all cylinders. Let Π be the group of all finite permutations of $\mathbb{N}$. If we let $\mathcal{A}$ be the set of all exchangeable events; that is, events $E \in \mathcal{G}$ for which $\pi^{-1}E = E$ for all permutations $\pi \in \Pi$. If we let $\mathcal{P} = \{P : \forall \pi \in \Pi,\ \forall E \in \mathcal{G},\ P(\pi^{-1}E) = P(E)\}$ be the set of all exchangeable probability measures, we know that $(\Omega, \mathcal{G}, \mathcal{P})$ is a Dynkin space with sufficient σ-algebra $\mathcal{A}$ and the set of extreme points $\mathcal{S}$ is given by the models $P \in \mathcal{P}$ that take on 0-1 values on the sufficient σ-algebra $\mathcal{A}$. For example, if $S = \{0, 1\}$, the set of structured models could be given by the iid Bernoulli distributions with success parameter in between $\underline{p} < \bar{p}$; that is, $\mathcal{P}_0 = \{q = \times_{n \in \mathbb{N}} \mathbf{p} : \mathbf{p} \sim Ber(p) \text{ for some } p \in (\underline{p}, \bar{p}).\}$

3. Subjective Rationality, Coherence, and Consistency

In all the following discussion, we fix a structured space $(\Omega, \mathcal{G}, \mathcal{M})$ with a best-fit map $\mathfrak{q}$ and sufficient σ-algebra $\mathcal{A}$ satisfying the properties outlined in Definition 1.

3.1. Subjective Rationality.

First of all, we assume that the preferences of the decision maker satisfy behavioral axioms capturing the idea of *subjective rationality.* We state some axioms characterizing the notion that the preferences of the decision maker are subjectively rational.

Axiom 1 (Subjective Rationality):

(i) Weak Order. $\succsim$ *is complete and transitive.*

(ii) Monotonicity. *For all* $f, f' \in \mathcal{F}$*, if* $f(\omega) \succsim f'(\omega)$ *for all* $\omega \in \Omega$*, then* $f \succsim f'$*.*

(iii) Mixture Continuity. *If $f, f', f'' \in \mathcal{F}$, the sets $\{\alpha \in [0,1] : \alpha f' + (1-\alpha) f'' \succsim f\}$ and $\{\alpha \in [0,1] : f \succsim \alpha f' + (1-\alpha) f''\}$ are both closed.*

(iv) Risk Independence. *For all $x, y, z \in X$ and $\alpha \in [0,1]$,*

$$x \succsim y \iff \alpha x + (1-\alpha) z \succsim \alpha y + (1-\alpha) z \ .$$

(v) Uncertainty Aversion. *For all $f, f' \in \mathcal{F}$ and $\alpha \in (0,1)$,*

$$f \sim f' \implies \alpha f' + (1-\alpha) f \succsim f \ .$$

(vi) Unboundedness. *There exist $x, y \in X$ such that $x \succ y$ and for all $\alpha \in (0,1)$, there are $z, z' \in X$ such that*

$$\alpha z + (1-\alpha) y \succ x \succ y \succ \alpha x + (1-\alpha) z' \ .$$

The first four requirements of subjective rationality guarantee that the preferences are a continuous and monotone weak order satisfying independence when restricted to constant acts. Then, the theorem of Herstein and Milnor (1953) implies that the preferences are represented on X by an affine utility u. If we interpret the mixture space X as the set of simple lotteries over outcomes, these axioms imply that the decision maker evaluates lotteries - i.e., constant acts that are not affected by ambiguity but only involve risk - according to their objective expected utility. Requirement (v) is the classical axiom capturing a preference for hedging due to Schmeidler (1989), and is usually interpreted in terms of averse attitudes towards uncertainty. The last requirement could be substituted with the much weaker non-triviality assumption that there exist constant acts $x, y \in X$ such that $x \succ y$. We go with this stronger requirement for technical convenience, as it guarantees that the utility over consequences u will be unbounded above and below.

Finally, the next axiom guarantees that preferences are robust to small perturbations and guarantees the countable additivity of the subjective probabilities.

Axiom 2 (Monotone Continuity): *For all* $f, f' \in \mathcal{F}$ *and* $x \in X$, *for all* $(E_n)_{n\in\mathbb{N}} \subseteq \mathcal{G}$ *such that* $E_1 \supseteq E_2 \supseteq \cdots$ *and* $\bigcap_{n\in\mathbb{N}} E_n = \emptyset$, *if* $f \succ f'$, *then, there exists* $n_0 \in \mathbb{N}$ *such that* $xE_{n_0}f \succ f'$.

3.2. **Coherence and Consistency.** Define for each structured model $m \in \mathcal{M}$ the set of states of the world for which the best-fit map would imply that m is the best approximation of the DGP. This is the set $E^m := \mathfrak{q}^{-1}(m)$ and notice that $E^m \in \mathcal{A}$. Moreover, given $p \in \Delta(\Omega)$, for each simple act f, define

$$\mathbb{E}_p[f] := \sum_{x\in \operatorname{Im} f} xp\left(f^{-1}(x)\right)$$

be the "average" of f according to the probability model p. Notice that since f has finite image and X is convex, $\mathbb{E}_p[f] \in X$ and it is the certainty equivalent of f for an Anscombe-Aumann EU maximizer who holds belief p over the state space Ω. The following axiom captures the idea that the preferences of the decision maker are coherent with the statistical framework embodied by the structured space.

Axiom 3 (Coherence):

(i) *For all structured models* $m \in \mathcal{M}$, E^m *is nonnull and* $fE^mh \succsim gE^mh$ *if and only if* $fE^mh' \succsim gE^mh'$ *for all* $f, g, h, h' \in \mathcal{F}$.

(ii) *For all* $m \in \mathcal{M}$ *and* $f, g, h \in \mathcal{F}$,

$$f = g \quad a.e.\ [m] \implies fE^mh \sim gE^mh\ .$$

(iii) *For all* $m \in \mathcal{M}$, *if* $p \ll m$ *but* $p \neq m$, *then there exist* $f \in \mathcal{F}$ *and* $x \in X$ *such that* $fE^mx \succsim x$ *but* $x \succ \mathbb{E}_p[f]$.

(iv) *For all* $x \in X$ *and* $f \in \mathcal{F}$, *the set* $\{m \in \mathcal{M} : fE^mx \succsim x\}$ *is measurable.*

Coherence requires that the decision maker's preferences are adapted to the statistical information implied by the structured models and the best-fit map. First of all, point *(i)* requires that for each structured model m, the preferences of the decision maker deem possible the event that m is the best approximation of the true DGP. We can think of this as a parsimony requirement: if the decision maker thought that a structured model could never be the best-fit one, then she might just as well drop it altogether. Moreover, the second part of the first point requires that the decision maker is able to identify for each structured model the event that such model is the best approximation of the DGP and make conditional assessments of the acts based on this event. In particular, this guarantees that we can define nontrivial preferences $\succsim^m$ conditional on a structured model $m \in \mathcal{M}$ being the best approximation to the true model in an unambiguous way: for all $f, g \in \mathcal{F}$,

$$f \succsim^m g \iff (\exists h \in \mathcal{F},\ fE^m h \succsim gE^m h) \ .$$

The second requirement ensures that the preferences of the decision maker recognize the substantive motivations underlying the selected structured models and incorporate the information provided by the best-fit map. Indeed, if two acts are equal with probability one according to a structured model $m \in \mathcal{M}$, the fact that the decision maker pays special attention to the model when it is the best-fit one is reflected by the fact that the two acts are ranked as indifferent conditional on the event that m is, indeed, the best approximation. The third point clarifies the interpretation of m being substantively motivated in the eyes of the decision maker compared to other models not in $\mathcal{M}$. This is reflected in the fact that for each non-structured model p that also assigns probability one to the event E^m, there exists a (possibly) uncertain act f that the decision maker would be willing to take over a deterministic outcome x conditional on m being the best-fit model even if x is strictly preferred to the EU certainty equivalent of f according to p. Finally, the last point is a measurability requirement of preferences with respect to the sufficient σ-algebra. To summarize, coherence implies that each structured model $m \in \mathcal{M}$ induces a well-defined and non-trivial conditional preference $\succsim^m$ that ranks as indifferent acts that are equal with probability one according to m and changes in a measurable fashion with respect to the structured models.

The next axiom is key in tying together the subjectively rational preferences with the set of structured models and the conditional preferences they induce.

Axiom 4 (Consistency): *For all* $f, f', g \in \mathcal{F}$,

$$(\forall m \in \mathcal{M}, \quad fE^m g \succsim f'E^m g) \implies f \succsim f' \ .$$

This assumption is analogous in nature to the consistency axiom introduced in Gilboa et al. (2010) and Cerreia-Vioglio et al. (2013). We can think of the set of structured models $\mathcal{M}$ as identifying an objective preference over acts. If an act f dominates act f' conditional on each structured model $m \in \mathcal{M}$, then f is objectively preferred by the decision maker to f'. Consistency requires that the subjectively rational preferences of the decision maker are informed by the objectively rational preferences.

3.3. **Misspecification Aversion.** We next state a conditional version of the axiom characterizing the variational preferences of Maccheroni et al. (2006). That is, the preferences after conditioning on the event that the structured model $m \in \mathcal{M}$ is the best-fit one satisfy a stronger form of independence, weak certainty independence, but they still do not need to satisfy full-fledged independence because of misspecification concerns.

Axiom 5 (Misspecification Aversion): *For all structured models* $m \in \mathcal{M}$*,* $f, f' \in \mathcal{F}$*,* $x, y \in X$*, and* $\alpha \in (0,1)$*,*

$$\alpha f + (1-\alpha)x \succsim \alpha f' E^m f + (1-\alpha)x \implies \alpha f + (1-\alpha)y \succsim \alpha f' E^m f + (1-\alpha)y \ .$$

We interpret Axiom 5 as capturing the idea that the decision maker is aware that the set of structured models is possibly misspecified and is concerned about it. Recall that we interpret ambiguity as the lack of information needed to pin down a unique probability distribution over states of the world. Now, suppose the decision maker was able to observe sufficient information to determine that a structured model m is the best-fit among all those in $\mathcal{M}$. If she was completely certain that the true DGP is included in $\mathcal{M}$, she should conclude as a matter of fact that m *is* the correct description of the uncertainty about the states. If such were the case, having received the missing information to determine the DGP, there is no reason why the decision maker's preferences should exhibit any ambiguity aversion but should instead behave according to the subjective expected utility tenets. The fact that even after being told that m is the best-fit structured model, the decision maker's preferences might still feature violations of independence implies that she does not trust that the best-fit model m is, in fact, the true DGP, reflecting a concern for the set of structured models being misspecified.

4. Representation of Misspecification Averse Preferences

In this section we discuss our main representation results. We begin by defining the preferences under analysis.

Definition 2 (Misspecification Averse Preferences): A preference relation $\succsim$ on $\mathcal{F}$ is said to be *Misspecification Averse* if it satisfies Axioms 1, 2, 3, 4, and 5.

As a first step, we provide a representation for the preferences $\succsim^m$ conditional on the event that the structured model $m \in \mathcal{M}$ is the best-fit one. The result is that each $\succsim^m$ is a variational preference (Maccheroni et al., 2006).

Proposition 1: *Suppose* $(\Omega, \mathcal{G}, \mathcal{M})$ *is a structured space and* $\succsim$ *is a misspecification averse preference relation. Then, there exist an affine and surjective utility function* $u : X \to \mathbb{R}$ *and a convex statistical distance*[11] $c : \Delta \times \mathcal{M} \to [0, \infty]$*, such that for each* $m \in \mathcal{M}$*,* $f, g \in \mathcal{F}$*,*

$$f \succsim^m g \iff I^m(u(f)) \geq I^m(u(g))$$

where $I^m : B(\mathcal{G}) \to \mathbb{R}$ *is defined as*

$$I^m(\varphi) = \min_{p \in \Delta} \left\{ \int_\Omega \varphi dp + c(p, m) \right\} \tag{5}$$

for all $\varphi \in B(\mathcal{G})$ *and satisfies for all* $\varphi, \varphi' \in B(\mathcal{G})$*,*

$$\varphi = \varphi' \text{ a.e. } [m] \implies I^m(\varphi) = I^m(\varphi').$$

Moreover, u *is unique up to positive affine transformations, and* c *is unique given* u*.*

We can interpret the result in Proposition 1 in terms of a robust approach to the possibility of misspecification. Suppose that the decision maker has observed sufficient information to determine that m is the best-fit structured model she has available. Because of the possibility of misspecification, in evaluating an act f conditional on this information, the decision maker forms a variational evaluation of the act f

$$V^m(f) \coloneqq I^m(u(f)) = \min_{p \in \Delta} \left\{ \int_\Omega u(f) dp + c(p, m) \right\} . \tag{6}$$

The statistical distance $c(\cdot, m)$ captures how distant in a statistical sense an unstructured model p is from the structured model m. In particular, since the decision maker

[11] See Appendix A for a rigorous definition of the notion of statistical distance.

is concerned that m might not be an accurate approximation of the true DGP, she also takes into account other models p that are not too far apart from m. The index $c(\cdot, m)$ captures exactly the decision maker's confidence in the structured model m. When $c(\cdot, m)$ is (uniformly) lower, the decision maker potentially takes into account a larger set of models around m in evaluating an act; this reflects a lower trust in m or, conversely, a higher aversion to misspecification, as reflected in the concern that m is not a good approximation even if she knows it is the best structured model. An important and tractable case is when the misspecification index takes the form $c(\cdot, m) = \lambda R(\cdot||m)$ for all structured models $m \in \mathcal{M}$, where $\lambda > 0$ is a parameter of misspecification aversion. In this case, the misspecification concern is proportional to the relative divergence with respect to the structured model, and it is uniform across structured models (see Lanzani (2022)). In this case, a higher aversion towards misspecification is captured by a lower parameter λ. We now make this intuition about the statistical distance $c(\cdot, m)$ precise by adapting to the present context the well-established notion of comparative uncertainty aversion due to Ghirardato and Marinacci (2002). Given two preferences $\succsim_1$ and $\succsim_2$, we say that $\succsim_1$ is *more misspecification averse* than $\succsim_2$ if for all $m \in \mathcal{M}$, $f \in \mathcal{F}$ and $x \in X$,

$$fE^m x \succsim_1 x \implies fE^m x \succsim_2 x \tag{7}$$

The idea behind this notion is that also in this case, constant acts are unaffected by the possibility that the set of structured models is misspecified, since they are non-stochastic and, therefore, their evaluation does not depend on the probabilistic assessment of state uncertainty. Therefore, if it is true that after conditioning on any given structured model $m \in \mathcal{M}$, a decision maker is not concerned enough about misspecification to choose a constant act over an uncertain one, a fortiori, that should also be true for a less misspecification averse decision maker. The following result shows that this definition agrees with the notion that the statistical distance in the representation is an index of misspecification aversion.

PROPOSITION 2: *Suppose* $(\Omega, \mathcal{G}, \mathcal{M})$ *is a structured space and* $\succsim_1$ *and* $\succsim_2$ *are two misspecification averse preference relations. Then,* $\succsim_1$ *is more misspecification averse than* $\succsim_2$ *if and only if* u_2 *is a positive affine transformation of* u_1 *and, after normalizing* $u_1 = u_2$, $c_1(\cdot, m) \leq c_2(\cdot, m)$ *for all structured models* $m \in \mathcal{M}$.

From a mathematical standpoint, note that each function I^m can be seen as a non-linear expectation with respect to the structured model $m \in \mathcal{M}$. Indeed, while it fails to be linear, it satisfies many other characteristic properties of expectations, like monotonicity, normalization and same evaluation of functions that are almost surely equal. The next corollary shows that by focusing on structured space, we are able to find a non-linear conditional expectation given $\mathcal{A}$ that is common to all structured models $m \in \mathcal{M}$.

COROLLARY 3: *Suppose* $(\Omega, \mathcal{G}, \mathcal{M})$ *is a structured space and* $(I^m)_{m\in\mathcal{M}}$ *are defined as in* (5). *Then, there exists a generalized common conditional expectation of* $(I^m)_{m\in\mathcal{M}}$ *given* $\mathcal{A}$. *This is a map* $I_{\mathcal{A}} : B(\mathcal{G}) \to \mathbb{R}^{\Omega}$ *such that for all* $\varphi \in B(\mathcal{G})$, $I_{\mathcal{A}}(\varphi)$ *is in* $B(\mathcal{A})$, $I_{\mathcal{A}}(\varphi)(\omega) = I^{\mathfrak{q}(\omega)}(\varphi)$ *for all* $\omega \in \Omega_0$ *and for all* $A \in \mathcal{A}$ *and* $m \in \mathcal{M}$,

$$I^m\left(I_{\mathcal{A}}(\varphi)\chi_A\right) = I^m(\varphi\chi_A).$$

Given the representation of the conditional preferences given in Proposition 1, we are able to associate to each act $f \in \mathcal{F}$ a function $m \mapsto I(f, m) := I^m(f)$ that maps each structured model m to the misspecification-robust evaluation of act f conditional on m being the best-fit model. The axiom of Consistency then implies that if $I(f, m) \geq I(g, m)$ for all $m \in \mathcal{M}$, then f should be preferred to g. That is, if for each structured model m, the corresponding misspecification-robust evaluation of an act f is always higher than the robust evaluation of another act g, the decision maker is confident that act f is better than act g, and her preferences will follow suit. As remarked in Cerreia-Vioglio et al. (2020), this exemplifies the special status of structured models over unstructured ones. If the misspecification-robust evaluations according to each structured model rank unanimously an act over another, this is sufficient for the decision maker to decide to pick the first one. However, in general, the set of structured models will not provide a unanimous robust ranking of every pair of acts. The first main result is that the representation of misspecification averse preferences will only depend on $I(f, \cdot)$; that is, there exists a monotone, continuous, and quasi-concave aggregator of these misspecification-robust evaluations that represents the preferences of the decision maker.

THEOREM 4: *Suppose* $(\Omega, \mathcal{G}, \mathcal{M})$ *is a structured space. The following are equivalent:*

(i) $\succsim$ *is a misspecification averse preference relation,*

(ii) there exist a surjective utility function $u : X \to \mathbb{R}$*, a convex statistical distance* $c : \Delta \times \mathcal{M} \to [0, \infty]$*, a monotone, normalized, quasiconcave, and lower semicontinuous functional* $\hat{I} : B(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$*, which is continuous on* $B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$ *such that for all* $f, g \in \mathcal{F}$*,*

$$f \succsim g \iff \hat{I}\left(I(u(f), \cdot)\right) \geq \hat{I}\left(I((u(g), \cdot)\right) \tag{8}$$

where for all $m \in \mathcal{M}$*,* $I(\cdot, m) = I^m(\varphi)$ *is given as in Proposition 1:*

$$\forall \varphi \in B(\mathcal{G}),\ I(\varphi, m) = I^m(\varphi) = \min_{p \in \Delta} \left\{ \int_\Omega \varphi dp + c(p, m) \right\}.$$

Moreover, u *is unique up to positive affine transformations, and* c *and* $\hat{I}$ *are unique given* u*.*

We already discussed how $c(\cdot, m)$ can be interpreted as an index of the decision maker's uncertainty aversion. On the other hand, quasiconcavity of $\hat{I}$ reflects the decision maker aversion towards uncertainty about the identity of the best-fit model, given that she lacks the sufficient information represented by the sigma-algebra $\mathcal{A}$ to pin down the best approximation of the DGP at her disposal. We now make precise the idea that $\hat{I}$ captures attitudes towards the uncertainty regarding the best-fit model. To this end, say that $\succsim_1$ is *more averse to model ambiguity* than $\succsim_2$ if for all $f \in \mathcal{F}(\mathcal{A})$ and $x \in X$,

$$f \succsim_1 x \implies f \succsim_2 x. \tag{9}$$

The intuition for this definition is that acts that are measurable with respect to the sufficient information $\mathcal{A}$ are exactly those acts that are only affected by the uncertainty regarding what is the best approximation among the set of structured models but not by misspecification concerns regarding any structured models (notice that they need to be constant on each event E^m). Therefore, the definition above states that if $\succsim_1$ is more averse to model ambiguity than $\succsim_2$ then, whenever model ambiguity considerations are not enough for the first decision maker to prefer the certain outcome x to the act f that is affected by ambiguity about the best-fit model, then definitely they should not be enough for the less averse decision maker. We have the following comparative statics result.

PROPOSITION 5: *Suppose $(\Omega, \mathcal{G}, \mathcal{M})$ is a structured space and $\succsim_1$ and $\succsim_2$ are two misspecification averse preference relations. Then, $\succsim_1$ is more model ambiguity averse than $\succsim_2$ if and only if u_2 is a positive affine transformation of u_1 and, after normalizing $u_1 = u_2$, $\hat{I}_1 \leq \hat{I}_2$.*

Since $\hat{I}_1$ and $\hat{I}_2$ are normalized, they can be interpreted as certainty equivalents of uncertain bets on the likelihood of which model is the best-fit one. The result can then be taken as stating that $\succsim_1$ is more averse to model ambiguity than $\succsim_2$ if DM 1 is willing to accept lower certainty equivalents than DM 2 as compensation for uncertain bets over the likelihood of the best approximation in $\mathcal{M}$. In this sense, Proposition 5 allows us to interpret the aggregator $\hat{I}$ as incorporating the decision maker's attitudes towards uncertainty about the identity of the best structured model. This result, together with Proposition 2, clarifies how representation (8) achieves a separation of attitudes regarding the ambiguity about the identity of the best-fit model and misspecification concerns. Indeed, aversion to model ambiguity is captured by the aggregator $\hat{I}$, while the statistical distance $c(\cdot, m)$ is an index of the degree of aversion to the possibility that the set of structured models is misspecified.

The abstract form of $\hat{I}$ in the general representation of Theorem 4 is due to the fact that no behavioral assumptions regarding the independence properties of the preference relation $\succsim$ have been made other than risk independence. The next two results characterize two specific shapes of the monotone aggregator of the misspecification-robust evaluations. The first result provides a foundation for a Bayesian version of the misspecification averse preferences, where the DM forms a subjective belief capturing her uncertainty regarding the identity of the best-fit structured model in $\mathcal{M}$.

Theorem 6: $(\Omega, \mathcal{G}, \mathcal{M})$ *is a structured space. The following are equivalent:*

(i) $\succsim$ *is a misspecification averse preference relation whose restriction to* $\mathcal{F}(\mathcal{A})$ *satisfies Savage (1954)'s Axioms P2-P6,*

(ii) there exist a surjective and affine utility function $u : X \to \mathbb{R}$, *a convex statistical distance* $c : \Delta \times \mathcal{M} \to [0, \infty]$, *a strictly increasing, continuous, and concave function* $\phi : \mathbb{R} \to \mathbb{R}$ *and a nonatomic prior* $\mu \in \Delta^{\sigma}(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$ *such that* $\succsim$ *is represented on* $\mathcal{F}$ *by:*

$$V(f) := \int_{\mathcal{M}} \phi \left(\min_{p \in \Delta} \int_{\Omega} u(f) dp + c(p, m) \right) d\mu(m) \ . \tag{10}$$

Moreover, u *is unique up to positive affine transformations,* c *is unique given* u, ϕ *is unique up to positive affine transformations given* u, *and* μ *is unique.*

As before, the decision maker's concern for misspecification is captured by the fact that even conditioning on the information to determine m as the structured model, she still takes into account models that are not structured but are close enough to m. In this case, the perception of uncertainty regarding the identity of the best-fit model in the absence of the information in the sufficient sigma-algebra $\mathcal{A}$ and the attitudes towards this uncertainty are captured, respectively, by the Bayesian prior μ over the set of structured models and the index of uncertainty aversion ϕ. The subjective belief μ quantifies what structured models the decision maker considers more likely to be good approximations of the true DGP. The nonlinearity of ϕ captures the negative attitude exhibited by the decision maker towards this ambiguity about the best-fit model. The Bayesian criterion (10) can be seen as an extension of the smooth ambiguity model of Klibanoff et al. (2005) to incorporate misspecification concerns. We can recover the smooth ambiguity model by letting the misspecification aversion index c go to infinity (except on the diagonal, where it is always 0). This is equivalent to taking a limit case where the decision maker is neutral to misspecification. As already remarked in the introduction, this criterion becomes the average robust control criterion axiomatized by Lanzani (2022) when the decision maker is neutral towards the ambiguity regarding the identity of the best-fit model. This would, indeed, imply that the index ϕ is affine. The relative entropy formulation of the misspecification aversion index $c(\cdot, m) = \lambda R(\cdot || m)$ could be obtained by imposing suitable versions of the multiplier preferences axioms discussed by Strzalecki (2011).

Finally, the next theorem shows that the criterion axiomatized in Cerreia-Vioglio et al. (2020) can arise as a special case of the representation in Theorem 4 when we assume that preferences exhibit a cautious attitude with respect to the uncertainty about the best-fit model.

Axiom 6 ($\mathcal{M}$-Caution): *For all $f \in \mathcal{F}$ and $x \in X$,*

$$\exists m \in \mathcal{M},\ x \succ fE^m x \implies x \succsim f.$$

This axiom is the conceptually similar to the caution axiom in Gilboa et al. (2010). Indeed, the set of structured models induces a typically incomplete dominance relation $\succsim_{\mathcal{M}}$, where for all $f, g \in \mathcal{F}$,

$$f \succsim_{\mathcal{M}} g \iff \forall m \in \mathcal{M},\ f \succsim^m g.$$

If $f \succsim_{\mathcal{M}} g$, this means that f is better than g according to each structured model $m \in \mathcal{M}$ after taking into account misspecification concerns. Because of the substantive motivation the decision maker attaches to the set of structured models, when $f \succsim_{\mathcal{M}} g$, the decision maker is sure that f is better than g. Then, Axiom 6 can be rewritten as the requirement that if $f \not\succsim_{\mathcal{M}} x$, then $x \succsim f$. The interpretation is that if the decision maker is not sure that the uncertain act f is better than the constant (and therefore unaffected by uncertainty considerations) act x, then she should behave cautiously and prefer the certain act over the uncertain one. We also impose the following technical axiom.

Axiom 7 ($\mathcal{M}$-Lower Semicontinuity): *For all $x \in X$ and $f \in \mathcal{F}$, the set $\{m \in \mathcal{M} : x \succsim fE^m x\}$ is closed.*

This axiom is a strengthening of requirement (*iii*) in the axiom of Coherence (it requires closedness and not only measurability) and it is only needed to ensure that minima are achieved in the criterion. The following result shows that $\mathcal{M}$-Caution delivers the criterion of Cerreia-Vioglio et al. (2020).

Theorem 7: *Suppose $(\Omega, \mathcal{G}, \mathcal{M})$ is a structured space and $\mathcal{M}$ is compact.*[12] *The following are equivalent:*

[12]As for Axiom 7, closedness of $\mathcal{M}$ is only needed to ensure that minima are achieved.

(i) $\succsim$ *is a misspecification averse preference relation satisfying Axioms 6 and 7,*

(ii) there exists a surjective utility function $u : X \to \mathbb{R}$*, a convex statistical distance* $c : \Delta \times \mathcal{M} \to [0, \infty]$ *such that* $\succsim$ *is represented on* $\mathcal{F}$ *by:*

$$V(f) = \min_{p \in \Delta} \int_{\Omega} u(f) dp + \min_{m \in \mathcal{M}} c(p, m). \tag{11}$$

Moreover, u is unique up to positive affine transformations, and c is unique given u.

Notice that $\min_{m \in \mathcal{M}} c(m', m) = 0$ for all structured models $m' \in \mathcal{M}$. Therefore, $C_{\mathcal{M}}(\cdot) := \min_{m \in \mathcal{M}} c(\cdot, m)$ can be seen as a statistical distance between probability distributions and the set of structured models $\mathcal{M}$ capturing the degree of misspecification concern of the decision maker, when she takes a worst-case scenario approach to the uncertainty regarding what is the best structured model.

5. Conclusion

This paper provides an axiomatic foundation of general preferences that are misspecification averse. We study a framework where the decision maker formulates a possibly misspecified set of structured models that she considers plausible descriptions of the environment. We introduce the notion of a best-fit map that identifies the most suitable approximation of the true DGP based on (in principle) observable states. This allows us to discern between the decision maker's concern about the set of structured models being misspecified and negative attitudes towards the uncertainty about what structured models are more likely to be the best description of the environment. The main result is that the decision maker's preferences are a monotone and quasiconcave aggregation of misspecification-robust evaluations based on each structured model. In particular, this representation achieves a separation of attitudes towards model ambiguity, captured by the aggregator, and misspecification concerns, captured by the misspecification-robust conditional evaluations. Many specific shapes of the aggregator can be obtained by imposing additional suitable behavioral axioms on the decision maker's preferences. We show that two important decision criteria recently introduced in the literature by Lanzani (2022) and Cerreia-Vioglio et al. (2020) fall within the general class of misspecification averse preferences we studied. In particular, we provide specific axioms to obtain the Bayesian aggregator and the cautious criterion from the general case.

References

Al-Najjar, N. I. and L. De Castro (2014): "Parametric representation of preferences," *Journal of Economic Theory*, 150, 642–667.

Aliprantis, C. and K. Border (2007): *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer.

Amarante, M. (2009): "Foundations of neo-Bayesian statistics," *Journal of Economic Theory*, 144, 2146–2173.

Ba, C. (2021): "Robust Model Misspecification and Paradigm Shifts," *arXiv preprint arXiv:2106.12727*.

Berk, R. H. (1966): "Limiting behavior of posterior distributions when the model is incorrect," *The Annals of Mathematical Statistics*, 37, 51–58.

Billingsley, P. (1995): *Probability and Measure*, Wiley Series in Probability and Statistics, Wiley.

Cerreia-Vioglio, S., L. P. Hansen, F. Maccheroni, and M. Marinacci (2020): "Making decisions under model misspecification," *arXiv preprint arXiv:2008.01071*.

Cerreia-Vioglio, S., F. Maccheroni, and M. Marinacci (2022): "Ambiguity aversion and wealth effects," *Journal of Economic Theory*, 199, 104898.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011): "Uncertainty averse preferences," *Journal of Economic Theory*, 146, 1275–1330.

——— (2013): "Ambiguity and robust statistics," *Journal of Economic Theory*, 148, 974–1049.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and A. Rustichini (2014): "Niveloids and their extensions: Risk measures on small domains," *Journal of Mathematical Analysis and Applications*, 413, 343–360.

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

Dynkin, E. (1978): "Sufficient statistics and extreme points," *The Annals of Probability*, 6, 705–730.

Epstein, L. G. and K. Seo (2010): "Symmetry of evidence without evidence of symmetry," *Theoretical Economics*, 5, 313–368.

Esponda, I. and D. Pouzo (2016): "Berk–Nash equilibrium: A framework for modeling agents with misspecified models," *Econometrica*, 84, 1093–1130.

Frick, M., R. Iijima, and Y. Ishii (2022): "Belief Convergence under Misspecified Learning: A Martingale Approach," *The Review of Economic Studies*.

Fudenberg, D. and G. Lanzani (2023): "Which misspecifications persist?" *Theoretical Economics*, 18, 1271–1315.

Fudenberg, D., G. Lanzani, and P. Strack (2021): "Limit points of endogenous misspecified learning," *Econometrica*, 89, 1065–1098.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004): "Differentiating ambiguity and ambiguity attitude," *Journal of Economic Theory*, 118, 133–173.

Ghirardato, P. and M. Marinacci (2002): "Ambiguity made precise: A comparative foundation," *Journal of Economic Theory*, 102, 251–289.

Gilboa, I., F. Maccheroni, M. Marinacci, and D. Schmeidler (2010): "Objective and subjective rationality in a multiple prior model," *Econometrica*, 78, 755–770.

Hansen, L. P. and T. J. Sargent (2001): "Robust control and model uncertainty," *American Economic Review*, 91, 60–66.

——— (2022): "Structured ambiguity and model misspecification," *Journal of Economic Theory*, 199, 105165.

He, K. and J. Libgober (2021): "Evolutionarily Stable (Mis) specifications: Theory and Applications," in *Proceedings of the 22nd ACM Conference on Economics and Computation*, 587–587.

Herstein, I. N. and J. Milnor (1953): "An axiomatic approach to measurable utility," *Econometrica, Journal of the Econometric Society*, 291–297.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849–1892.

Klibanoff, P., S. Mukerji, and K. Seo (2014): "Perceived ambiguity and relevant measures," *Econometrica*, 82, 1945–1978.

Lanzani, G. (2022): "Dynamic Concern for Misspecification," .

Maccheroni, F., M. Marinacci, and A. Rustichini (2006): "Ambiguity aversion, robustness, and the variational representation of preferences," *Econometrica*, 74, 1447–1498.

Mauldin, R. D., D. Preiss, and H. v. Weizsacker (1983): "Orthogonal transition kernels," *The Annals of Probability*, 970–988.

Mu, X., L. Pomatto, P. Strack, and O. Tamuz (2021): "Monotone additive statistics," *arXiv preprint arXiv:2102.00618*.

Savage, L. J. (1954): *The foundations of statistics*, New York: John Wiley & Sons.

Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.

Strzalecki, T. (2011): "Axiomatic foundations of multiplier preferences," *Econometrica*, 79, 47–73.

Wald, A. (1950): *Statistical decision functions*, New York: Wiley.

Weis, L. W. (1984): "A characterization of orthogonal transition kernels," *The Annals of Probability*, 1224–1227.

Appendix

Appendix A. Mathematical Preliminaries

A.1. **Basic Notions.** Given an arbitrary measurable space $(Y,\mathcal{Y})$, we denote by $\Delta(Y,\mathcal{Y})$ and $\Delta^\sigma(Y,\mathcal{Y})$ respectively the space of finitely and countably additive probability measures on $(Y,\mathcal{Y})$. Sometimes, we will omit making explicit reference to the σ-algebra whenever no ambiguities can arise. Since both these spaces can be identified with subsets of the dual space of $B_0(Y,\mathcal{Y})$, the space of $\mathcal{Y}$-measurable simple functionals mapping Y to the real line, endowed with the supnorm $||\cdot||_\infty$, we endow them with the weak* topology. We endow $\Delta^\sigma(Y,\mathcal{Y})$ with the Borel σ-algebra generated by this topology; which is the same as the natural σ-algebra $\mathcal{D}^{Y,\mathcal{Y}}$ generated by the family of evaluations maps:

$$\forall E \in \mathcal{Y}, \quad E^* : \Delta^\sigma(Y,\mathcal{Y}) \to \mathbb{R},\ p \mapsto p(E)\ .$$

and any subset $\mathcal{Q}$ of Δ^σ, with the relative σ-algebra $\mathcal{D}^{Y,\mathcal{Y}}_{\mathcal{M}} := \mathcal{D}^{Y,\mathcal{Y}} \cap \mathcal{M}$. Moreover, denote by $B(Y,\mathcal{Y})$ the set of bounded $\mathcal{Y}$-measurable functionals from Y to $\mathbb{R}$. We know that $B(Y,\mathcal{Y})$ is the supnorm closure of $B_0(Y,\mathcal{Y})$.

Given a nonempty subset $\tilde{B}$ of $B(Y,\mathcal{Y})$, a functional $\Psi : \tilde{B} \to \mathbb{R}$ is said to be a *niveloid* if for all $\varphi, \varphi', \in \tilde{B}$,

$$\Psi(\varphi) - \Psi(\varphi') \leq \sup(\varphi - \varphi')$$

A niveloid is Lipschitz continuous with respect to the supnorm. Indeed:

$$\Psi(\varphi) - \Psi(\varphi') \leq \sup(\varphi - \varphi') \leq |\sup(\varphi - \varphi')| \leq \sup|\varphi - \varphi'| = ||\varphi - \varphi'||_\infty$$

$$\Psi(\varphi') - \Psi(\varphi) \leq \sup(\varphi' - \varphi) \leq |\sup(\varphi' - \varphi)| \leq \sup|\varphi' - \varphi'| = ||\varphi - \varphi'||_\infty$$

so that $|\Psi(\varphi) - \Psi(\varphi')| \leq ||\varphi - \varphi'||_\infty$ for all $\varphi, \varphi' \in \tilde{B}$. Moreover, the functional Ψ is said to be *normalized* if $\Psi(k) = k$ for all $k \in \mathbb{R}$ such that $k \in \tilde{B}$, where we identify each real number with the constant function yielding it everywhere. Finally, the functional Ψ is said to be *monotone* if whenever $\varphi, \varphi' \in \tilde{B}$ and $\varphi \geq \varphi'$, then $\Psi(\varphi) \geq \Psi(\varphi')$.[13] We say that Ψ is *monotone continuous* if for all $\varphi, \varphi' \in \tilde{B}$ and $k \in \tilde{B}$, for all monotone sequences $(E_n)_n \in \mathcal{Y}$ such that $E_n \downarrow \emptyset$, if $\Psi(\varphi) > \Psi(\varphi')$, then there exists $n_0 \in \mathbb{N}$ such that $\Psi(k\chi_{E_{n_0}} + \varphi\chi_{E^c_{n_0}})\varphi > \Psi(\varphi')$.

We define on $B(Y, \mathcal{Y})$ the *lattice operations* $\vee$ and $\wedge$ as follows: for all $\varphi, \varphi' \in B(Y, \mathcal{Y})$, $(\varphi \vee \varphi')(\omega) = \max\{\varphi(y), \varphi'(y)\}$ and $(\varphi \wedge \varphi')(\omega) = \min\{\varphi(y), \varphi'(y)\}$ for all $y \in Y$. We say that a nonempty subset L of $\subseteq B(Y, \mathcal{Y})$ is a *lattice* if for all $\varphi, \varphi' \in L$, $\varphi \vee \varphi', \varphi \wedge \varphi' \in L$. If $(\varphi_n)_N$ is a sequence of functions in $\subseteq B(Y, \mathcal{Y})$ and $\varphi \in B(Y, \mathcal{Y})$, we write $\varphi_n \to \varphi$ to mean that $(\varphi_n)_n$ converges uniformly to φ. If we want to stress that the uniformly convergent sequence is monotone, we write $\varphi_n \nearrow \varphi$ if $\varphi_n \leq \varphi_{n+1}$ for all $n \in \mathbb{N}$ and $\varphi_n \searrow \varphi$ if $\varphi_n \geq \varphi_{n+1}$ for all $n \in \mathbb{N}$. Finally, we write $\varphi_n \uparrow \varphi$ if $\varphi_n \leq \varphi_{n+1}$ for all $n \in \mathbb{N}$ and $(\varphi_n)_n$ converges pointwise to φ and, similarly, $\varphi_n \downarrow \varphi$ if $\varphi_n \geq \varphi_{n+1}$ for all $n \in \mathbb{N}$ and $(\varphi_n)_n$ converges pointwise to φ.

A.2. **Probabilities and Statistical Distances.** We now discuss some basic mathematical notions about probabilities and statistical distances. Fix an arbitrary measurable space $(Y, \mathcal{Y})$. For any $p, q \in \Delta(Y, \mathcal{Y})$, we write $p \ll q$ to denote that p is *absolutely continuous* with respect to q. Moreover, if $q \in \Delta(Y, \mathcal{Y})$ and f and g are $\mathcal{Y}$-measurable functions mapping Y to some arbitrary set, we write $f = g$ *a.e.* $[q]$ whenever $q(\{y \in Y : f(y) = g(y)\}) = 1$. As it is standard in measure-theoretic contexts, we assume throughout the convention $0 \cdot \infty = 0$. If f is a function mapping Y to some measurable space, we denote by $\sigma(f)$ the σ-algebra generated by f.

[13]See Maccheroni et al. (2006) and Cerreia-Vioglio et al. (2014) for an in-depth discussion of niveloids and their properties.

Given a convex subset C of $\Delta(Y, \mathcal{Y})$ and an extended real valued function $\varphi : C \to \bar{\mathbb{R}}$, we denote by $\operatorname{dom}\varphi$ the effective domain of φ, that is the subset of its domain on which φ takes on finite values; that is, $\operatorname{dom}\varphi \coloneqq \{p \in C : |\varphi(p)| < \infty\}$. Moreover, we say such function φ to be *grounded* if $\inf_{p\in C} \varphi(p) = 0$. Fix a subset $\mathcal{Q} \subseteq \Delta^\sigma(Y, \mathcal{Y})$ of countably additive probability measures. A function $c : \Delta(Y, \mathcal{Y}) \times \mathcal{Q} \to [0, \infty]$ is said to be a *statistical distance* if it satisfies the following two properties:

(i) for each $q \in \mathcal{M}$, $p = q$ implies $c(p, q) = 0$,

(ii) $c(\cdot, q)$ is lower semicontinuous for all $q \in \mathcal{Q}$.

Furthermore, a statistical distance c is convex if the section $c(\cdot, q)$ is a convex function for each $q \in \mathcal{Q}$ and is said to be a *divergence* if for all $q \in \mathcal{Q}$, $p \in \operatorname{dom} c(\cdot, q)$ implies that $p \ll q$.

Appendix B. Structured Spaces

Fix a measurable space $(\Omega, \mathcal{G})$ where $\mathcal{G}$ is a countably generated σ-algebra and a set of structured models $\mathcal{M} \subseteq \Delta^\sigma(\mathcal{G}) \coloneqq \Delta^\sigma(\Omega, \mathcal{G})$, where we denote by Let $\mathcal{D} \coloneqq \mathcal{D}^{\Omega,\mathcal{G}}$ and $\mathcal{D})\mathcal{M} \coloneqq \mathcal{D}^{\Omega,\mathcal{G}}_{\mathcal{M}}$ respectively the natural σ-algebra on $\Delta^\sigma(\mathcal{G})$ and the relative σ-algebra on $\mathcal{M}$. Throughout the section, assume that $(\Omega, \mathcal{G}, \mathcal{M})$ is a structured space with sufficient sub σ-algebra $\mathcal{A} \subseteq \mathcal{G}$ and best-fit map $\mathfrak{q} : \Omega \to \Delta^\sigma(\mathcal{G})$ satisfying the properties in Definition 1. In particular, recall that $E^m \coloneqq \mathfrak{q}^{-1}(m)$ and $m(E^m) = 1$ for all $m \in \mathcal{M}$. Denote by Λ the set of all the events in $\mathcal{G}$ that have probability either 0 or 1 according to all models $m \in \mathcal{M}$:

$$\Lambda \coloneqq \{E \in \mathcal{G} : \forall m \in \mathcal{M}, m(E) = 1 \text{ or } m(E) = 0\}.$$

Lemma B.1: *The σ-algebra generated by $\mathfrak{q}$ is in Λ: $\mathcal{A} = \sigma(\mathfrak{q}) \subseteq \Lambda$. In particular, $m(E) \in \{0, 1\}$ for all $E \in \mathcal{A}$ and structured model $m \in \mathcal{M}$.*

Proof of Lemma B.1: By definition of the σ-algebra $\mathcal{D}$, $\sigma(\mathfrak{q})$ is generated by the class:

$$\mathcal{C} := \left\{\mathfrak{q}^{-1}\left(\{p \in \Delta^\sigma(\mathcal{G}) : p(E) \leq x\}\right) : x \in [0,1],\ E \in \mathcal{G}\right\} .$$

Then, take any $x \in [0,1]$ and $E \in \mathcal{G}$. We have that for any $m \in \mathcal{M}$,

$$\begin{aligned} m\left(\mathfrak{q}^{-1}\left(\{p \in \Delta^\sigma(\mathcal{G}) : p(E) \leq x\}\right)\right) &= m\left(\{\omega \in \Omega : \mathfrak{q}^\omega(E) \leq x\}\right) \\ &= m\left(\{\omega \in \Omega : \mathfrak{q}^\omega(E) \leq x\} \cap E^m\right) \\ &= \begin{cases} 1 & \text{if } m(E) \leq x \\ 0 & \text{if } m(E) > 0 \end{cases}, \end{aligned}$$

and, therefore, $\mathfrak{q}^{-1}\left(\{p \in \Delta^\sigma(\mathcal{G}) : p(E) \leq x\}\right) \in \Lambda$, showing that $\mathcal{C} \subseteq \Lambda$.

It is clear that $\Omega, \emptyset \in \Lambda$ and that if $E \in \Lambda$, then $\Omega \setminus E \in \Lambda$. Moreover, if we take $(E)_{n\in\mathbb{N}} \subseteq \Lambda$, for each $m \in \mathcal{M}$, we have either of two cases. If $m(E_n) = 0$ for all $n \in \mathbb{N}$, then:

$$m(\cup_{n\in\mathbb{N}} E_n) \leq \sum_{n\in\mathbb{N}} m(E_n) = 0 \implies m(\cup_{n\in\mathbb{N}} E_n) = 0.$$

If, instead, there exists $k \in \mathbb{N}$ such that $m(E_k) = 1$, then:

$$m(\cup_{n\in\mathbb{N}} E_n) \geq m(E_k) = 1 \implies m(\cup_{n\in\mathbb{N}} E_n) = 1.$$

It follows that $\cup_{n\in\mathbb{N}} E_n \in \Lambda$. We can, thus, conclude that Λ is a σ-algebra containing $\mathcal{M}$ and, therefore, $\sigma(\mathfrak{q}) = \sigma(\mathcal{C}) \subseteq \Lambda$. ■

Suppose that $u : X \to \mathbb{R}$ is an affine and surjective function. If $\mathcal{E}$ is a sub-σ-algebra of $\mathcal{G}$, we can define the operator $u : \mathcal{F}(\mathcal{E}) \to B_0(\mathcal{E})$ as follows: for each $f \in \mathcal{F}(\mathcal{E})$,

$$u(f)(\omega) = u(f(\omega))$$

for all $\omega \in \Omega$.

Lemma B.2: *Suppose u is affine and surjective. Then, $u : \mathcal{F}(\mathcal{E}) \to B_0(\mathcal{E})$ is an affine operator and $\{u(f) : f \in \mathcal{F}(\mathcal{E})\} = B_0(\mathcal{E})$.*

Proof: Take any $f \in \mathcal{F}(\mathcal{E})$. Then, there exists a finite, measurable partition of Ω, $(E_i)_{i=1}^k \subseteq \mathcal{E}$, and consequences $(x_i)_{i=1}^k \subseteq X$ such that $f = \sum_{i=1}^k \chi_{E_i} x_i$. Then, for all E_i and for all $\omega \in E_i$,

$$u(f)(\omega) = u(f(\omega)) = u(x_i)$$

and therefore, $u(f) = \sum_{i=1}^k \chi_{E_i} u(x_i)$. Therefore, $u(f) \in B_0(\mathcal{E})$ for all $f \in \mathcal{F}(\mathcal{E})$ so that the operator is well-defined and $\{u(f) : f \in \mathcal{F}(\mathcal{E})\} \subseteq B_0(\mathcal{E})$. Moreover, take $\alpha \in (0,1)$ and $f, f' \in \mathcal{F}(\mathcal{E})$. We have that for all $\omega \in \Omega$,

$$\begin{aligned} u(\alpha f + (1-\alpha) f')(\omega) &= u((\alpha f(\omega) + (1-\alpha) f'(\omega)) \\ &= \alpha u(f(\omega)) + (1-\alpha) u(f'(\omega)) \\ &= \alpha u(f)(\omega) + (1-\alpha) u(f')(\omega) \end{aligned}$$

proving affinity. Finally, take any $\varphi \in B_0(\mathcal{E})$. Then, there exist a finite, measurable partition of Ω, $(E_i)_{i=1}^k \subseteq \mathcal{E}$, and reals $(r_i)_{i=1}^k \subseteq \mathbb{R}$ such that $\varphi = \sum_{i=1}^k \chi_{E_i} r_i$. Since $\operatorname{Im} u = \mathbb{R}$, for each r_i we can pick $x_i \in X$ such that $r_i = u(x_i)$. Setting $f = \sum_{i=1}^k \chi_{E_i} x_i$ we can see that $\varphi = u(f)$ and $\varphi \in \mathcal{F}(\mathcal{E})$. This shows that $B_0(\mathcal{E}) \subseteq \{u(f) : f \in \mathcal{F}(\mathcal{E})\}$. ■

Appendix C. Proof of Proposition 1

We say that a binary relation $\succsim$ over $\mathcal{F}$ is *solvable* if, for each act $f \in \mathcal{F}$, there exists a constant act $x_f \in X$ such that $x_f \sim f$. We call such (possibly non-unique) act the *certainty equivalent* of f. Next, we show that a preference relation that satisfies Axiom 1 is solvable.

Lemma B.3: *Suppose that $\succsim$ is a preference relation on $\mathcal{F}$ satisfying Axiom 1. Then, $\succsim$ is solvable.*

Proof of Lemma B.3: Fix any $f \in \mathcal{F}$. Since f takes on only finitely many values, we can pick x^* and x_* in X such that for all $\omega \in \Omega$, $x^* \succsim f(\omega) \succsim x_*$. By Axiom 1.ii, this implies that $x^* \succsim f \succsim x_*$. Now, $\{\alpha \in [0,1] : \alpha x^* + (1-\alpha)x_* \succsim f\}$ and $\{\alpha \in [0,1] : f \succsim \alpha x^* + (1-\alpha)x_*\}$ are closed by mixture continuity and are non-empty, since the first one contains 1 and the second one contains 0. Moreover, by completeness of $\succsim$, their union is the whole $[0,1]$. Since the closed, unit interval is connected, such sets must have a non-empty intersection. This shows the existence of $x_f \in X$ such that $x_f \sim f$. ■

We proceed by defining the preferences conditional on a given structured model $m \in \mathcal{M}$ being the best-fit model and show that they inherit some properties from the unconditional preferences. Let us first recall the following axioms characterizing the variational preferences axiomatized by Maccheroni et al. (2006).

Axiom B.1 (Variational):

- Weak Certainty Independence. *For all $f, f' \in \mathcal{F}$, $x, y \in X$, and $\alpha \in (0,1)$,*

$$\alpha f + (1-\alpha)x \succsim \alpha f' + (1-\alpha)x \implies \alpha f + (1-\alpha)y \succsim \alpha f' + (1-\alpha)y \ .$$

- Uncertainty Aversion. *For all $f, f' \in \mathcal{F}$ and $\alpha \in (0,1)$,*

$$f \sim f' \implies \alpha f' + (1-\alpha)f \succsim f \ .$$

Lemma B.4: *Suppose that $(\Omega, \mathcal{G}, \mathcal{M})$ is a structured space and that the preference relation $\succsim$ satisfies Axioms 1, 2, 3, 4, and 5. For all $m \in \mathcal{M}$, define $\succsim^m$ as follows: for all $f, f' \in \mathcal{F}$,*

$$f \succsim^m f' \iff \exists g \in \mathcal{F},\ fE^m g \succsim f'E^m g.$$

Then, $\succsim^m$ is well-defined, satisfies Axiom 1, 2, and B.1 and coincides with $\succsim$ when restricted to constant acts in X.

Proof of Lemma B.4: Fix any $m \in \mathcal{M}$ and consider $\succsim^m$ as defined in Equation B.4. We show that this is a well-defined binary relation over $\mathcal{F}$. Indeed, suppose that for $f, f' \in \mathcal{F}$, there exists some $g \in \mathcal{F}$ such that $fE^m g \succsim f'E^m g$. Then, Axiom 3 implies that $fE^m h \succsim f'E^m h$ for all $h \in \mathcal{F}$. Therefore, in the following, we just fix a $g \in \mathcal{F}$ and notice that $f \succsim^m f' \iff fE^m g \succsim f'E^m g$. Moreover, note that for any $f, f', g \in \mathcal{F}$ and $\alpha \in [0,1]$, $(\alpha f + (1-\alpha)f')E^m g = \alpha(fE^m g) + (1-\alpha)(f'E^m g)$. Indeed, if $\omega \in E^m$:

$$
\begin{aligned}
\left((\alpha f + (1-\alpha)f')E^m g\right)(\omega) &= (\alpha f + (1-\alpha)f')(\omega) \\
&= \alpha f(\omega) + (1-\alpha)f'(\omega) \\
&= \alpha(fE^m g)(\omega) + (1-\alpha)(f'E^m g)(\omega) \\
&= \left(\alpha(fE^m g) + (1-\alpha)(f'E^m g)\right)(\omega)
\end{aligned}
$$

and, if $\omega \in \Omega \setminus E^m$:

$$
\begin{aligned}
\left((\alpha f + (1-\alpha)f')E^m g\right)(\omega) &= g(\omega) \\
&= \alpha g(\omega) + (1-\alpha)g(\omega) \\
&= \alpha(fE^m g)(\omega) + (1-\alpha)(f'E^m g)(\omega) \\
&= \left(\alpha(fE^m g) + (1-\alpha)(f'E^m g)\right)(\omega)\ .
\end{aligned}
$$

Step 1: Weak Order. Take any $f, f' \in \mathcal{F}$. Then, since $\succsim$ is complete, it follows that either $fE^m g \succsim f'E^m g$ or $f'E^m g \succsim fE^m g$. That is, either $f \succsim^m f'$ or $f' \succsim^m f$, showing that $\succsim^m$ is complete. Moreover, suppose that there are $f, f', f'' \in \mathcal{F}$ such that $f \succsim^m f'$ and $f' \succsim^m f''$. Then, $fE^m g \succsim f'E^m g$ and $f'E^m g \succsim f''E^m g$. Since $\succsim$ is transitive, it follows that $fE^m g \succsim f'' E^m g$ and, therefore, that $f \succsim^m f''$. This shows that $\succsim^m$ is also transitive.

Step 2: Mixture Continuity. Take any $f, f', f'' \in \mathcal{F}$. We show that $\{\alpha \in [0,1] : \alpha f' + (1-\alpha) f'' \succsim^m f\}$ is closed. Indeed, take any $\alpha_0 \in [0,1]$ and let $g = f$:

$$
\begin{aligned}
& \alpha_0 \in \{\alpha \in [0,1] : \alpha f' + (1-\alpha) f'' \succsim^m f\} \\
\iff & \alpha_0 f' + (1-\alpha_0) f'' \succsim^m f \\
\iff & (\alpha_0 f' + (1-\alpha_0) f'') E^m g \succsim f E^m g \\
\iff & \alpha_0 (f' E^m f) + (1-\alpha_0)(f'' E^m f) \succsim f \\
\iff & \alpha_0 \in \{\alpha \in [0,1] : \alpha (f' E^m f) + (1-\alpha)(f'' E^m f) \succsim f\}
\end{aligned}
$$

so that $\{\alpha \in [0,1] : \alpha f' + (1-\alpha) f'' \succsim^m f\} = \{\alpha \in [0,1] : \alpha(f' E^m f) + (1-\alpha)(f'' E^m f) \succsim f\}$ and the latter is closed by Axiom 1. By an analogous argument, it follows that also $\{\alpha \in [0,1] : f \succsim^m \alpha f' + (1-\alpha) f''\}$ is closed. Hence, $\succsim^m$ satisfies mixture continuity.

Step 3: Weak Certainty Independence. Take any $f, f' \in \mathcal{F}$, $x, y \in X$ and $\alpha \in (0,1)$. Then,

$$
\alpha f + (1-\alpha) x \succsim^m \alpha f' + (1-\alpha) x \implies [\alpha f + (1-\alpha) x] E^m g \succsim [\alpha f' + (1-\alpha) x] E^m g
$$

and letting $g = \alpha f + (1-\alpha) x$ this implies that:

$$
\begin{aligned}
\alpha f + (1-\alpha) x &\succsim [\alpha f' + (1-\alpha) x] E^m [\alpha f + (1-\alpha) x] \\
&= \alpha f' E^m f + (1-\alpha) x.
\end{aligned}
$$

But, then, by Axiom 5,

$$\alpha f + (1-\alpha)y \succsim \alpha f' E^m f + (1-\alpha)y$$
$$= [\alpha f' + (1-\alpha)y] E^m [\alpha f + (1-\alpha)y],$$

which, then, implies that $\alpha f+(1-\alpha)y \succsim^m \alpha f'+(1-\alpha)y$. If follows that $\succsim^m$ satisfies Weak Certainty Independence. A fortiori, it satisfies Risk Independence.

Step 4: Non-triviality. Since E^m is nonnull, there must exist $f, f', g \in \mathcal{F}$ such that $fE^m g \succ f'E^m g$. Since f and f' are finite-valued, we can pick $x, y \in X$ so that $x \succsim f(\omega)$ and $f'(\omega) \succsim y$ for all $\omega \in E^m$. But then, monotonicity implies that

$$xE^m g \succsim fE^m g \succ f'E^m g \succsim yE^m g$$

and, by transitivity, $xE^m g \succ yE^m g$ so that $x \succ^m y$. It follows that $\succsim^m$ is non-trivial.

Step 5. $\succsim^m|_X = \succsim_X$. By Axiom 1, $\succsim$ is a non-trivial weak order satisfying mixture continuity and independence when restricted to X. By Steps 1-4, the same is true for $\succsim^m$. Then, by Herstein and Milnor (1953), there exist affine functions $u, u_m : X \to \mathbb{R}$ such that u represents $\succsim|_X$ and u_m represents $\succsim^m|_X$. Moreover, since both $\succsim$ and $\succsim^m$ are non-trivial, u and u_m are non-constant. Now, take any $x, y \in X$ such that $x \succsim y$. Then, for all $\omega \in \Omega$,

$$\omega \in E^m \implies (xE^m g)(\omega) = x \succsim y = (yE^m g)(\omega)$$
$$\omega \in \Omega \setminus E^m \implies (xE^m g)(\omega) = g(\omega) = (yE^m g)(\omega)$$

so that, since $\succsim$ satisfies reflexivity and monotonicity by Axiom 1, $xE^m g \succsim yE^m g$ and, therefore, $x \succsim^m y$. Thus, for all $x, y \in X$:

$$\begin{aligned} u(x) \geq u(y) &\implies x \succsim|_X y \\ &\implies x \succsim y \\ &\implies x \succsim^m y \\ &\implies x \succsim^m|_X y \\ &\implies u_m(x) \geq u_m(y) \ . \end{aligned}$$

By Corollary B.3 in Ghirardato et al. (2004), there exists $a \in \mathbb{R}_{++}$ abd $b \in \mathbb{R}$ such that $u = au_m + b$. This implies the claim.

Step 6: Monotonicity. Take $f, f' \in \mathcal{F}$ and assume that $f(\omega) \succsim^m f'(\omega)$ for all $\omega \in \Omega$. Since by Step 4, $\succsim^m|_X = \succsim_X$, it is also the case that $f(\omega) \succsim f'(\omega)$ for all $\omega \in \Omega$. Then, since $\succsim$ satisfies Axiom 1, reflexivity and monotonicity imply that $fE^m g \succsim f'E^m g$ and, therefore, $f \succsim^m f'$, proving the statement.

Step 7: Unboundedness. This follows immediately by Step 5.

Step 8. Uncertainty Aversion

Take any $f, f' \in \mathcal{F}$ and $\alpha \in (0, 1)$ and suppose that $f \sim^m f'$. Then, taking $g = f$ in the definition of $\succsim^m$ and since $\succsim$ satisfies Axiom 5, we have

$$\begin{aligned} f \sim^m f' &\implies f \sim f'E^m f \\ &\implies \alpha f + (1-\alpha) f'E^m f \succsim f \\ &\implies [\alpha f + (1-\alpha) f']E^m f \succsim fE^m f \\ &\implies \alpha f + (1-\alpha) f' \succsim^m f \end{aligned}$$

showing that $\succsim^m$ satisfies Uncertainty Aversion.

Step 9: Monotone Continuity.

Take any $f, f' \in \mathcal{F}$ such that $f \succ^m f'$, $x \in X$, and $(A_n)_{n\in\mathbb{N}} \subseteq \mathcal{G}$ such that $A_1 \supseteq A_2 \supseteq \cdots$ and $\bigcap_{n\in\mathbb{N}} A_n = \emptyset$. Taking $g = f$ in the definition of $\succsim^m$, we have that $f \succ f'E^m f$. Moreover, for each $n \in \mathbb{N}$, let $E_n := A_n \cap E^m$ and observe that $E_n = A_n \cap E^m \supseteq A_{n+1} \cap E^m = E_{n+1}$ and

$$\bigcap_{n\in\mathbb{N}} E_n = \bigcap_{n\in\mathbb{N}} (A_n \cap E^m) = (\bigcap_{n\in\mathbb{N}} A_n) \cap E^m = \emptyset \cap E^m = \emptyset.$$

Since $\succsim$ satisfies Axiom 2, we can find $n_0 \in \mathbb{N}$ such that $xE_{n_0}f \succ f'E^m f$. Moreover,

$$\begin{aligned} \omega \in E_{n_0} = A_{n_0} \cap E^m &\implies (xE_{n_0}f)(\omega) = x = ((xA_{n_0}f)E^m f)(\omega), \\ \omega \in E^m \setminus A_{n_0} &\implies (xE_{n_0}f)(\omega) = f(\omega) = ((xA_{n_0}f)E^m f)(\omega), \\ \omega \notin E^m &\implies (xE_{n_0}f)(\omega) = f(\omega) = ((xA_{n_0}f)E^m f)(\omega). \end{aligned}$$

Therefore, $(xA_{n_0}f)E^m f = xE_{n_0}f \succ f'E^m f$ which implies that $xA_{n_0}f \succ^m f'$ as we wanted to show. ■

We are now ready to prove Proposition 1.

Proof of Proposition 1: *(i) implies (ii)* Suppose that $(\Omega, \mathcal{G}, \mathcal{M})$ is a structured space and the preference relation $\succsim$ satisfies Axioms 1, 2, 3, 4, and 5. Since $\succsim$ is a non-trivial, continuous weak order satisfying independence when restricted to constant acts, we know by Herstein and Milnor (1953) that there exists an affine and non-constant function $u : X \to \mathbb{R}$ representing $\succsim$ over X. Moreover, such u is cardinally unique. Next, we show that $\operatorname{Im} u = \mathbb{R}$. Clearly, being u affine and X convex, $\operatorname{Im} u$ must be an interval. Pick $x, y \in X$ such that $x \succsim y$ and a monotonically decreasing sequence $(\alpha_n)_n \subseteq [0,1]$ such that $\alpha_n \to 0$. Then, by unboundedness, for each $n \in \mathbb{N}$, there exists $z_n, z'_n \in X$ such that:

$$\alpha_n z_n + (1-\alpha_n)y \succ x \succ y \succ \alpha_n z'_n + (1-\alpha_n)x$$

Since u represents $\succsim$ on X and is affine, this implies:

$$\alpha_n u(z_n) + (1-\alpha_n)u(y) > u(x) > u(y) > \alpha_n u(z'_n) + (1-\alpha_n)u(x)$$

and, rearranging:

$$u(z_n) > \frac{u(x)-u(y)}{\alpha_n} + u(y) \quad \text{and} \quad u(z'_n) < -\frac{u(x)-u(y)}{\alpha_n} + u(x)$$

for all $n \in \mathbb{N}$. Therefore, $(u(z_n))_n$ and $(u(z'_n))_n$ are sequences in $\operatorname{Im} u$, the first monotonically increasing and diverging to $+\infty$, the second monotonically decreasing and diverging to $-\infty$. This implies that $\operatorname{Im} u = \mathbb{R}$.

Now, fix $m \in \mathcal{M}$. By Lemma B.4, $\succsim^m|_X = \succsim|_X$. Therefore, $\succsim^m$ is represented by u when restricted to constant acts in X. Define the functional $I_0^m : B_0(\mathcal{G}) \to \mathbb{R}$ as follows: for each $\varphi \in B_0(\mathcal{G})$, $I^m(\varphi) = u(x_{f_\varphi})$ where $f_\varphi \in \mathcal{F}$ is chosen such that $\varphi = u(f_\varphi)$ and $x_{f_\varphi} \sim^m f_\varphi$. This functional is well-defined by Lemmas B.4 and B.2. Moreover, define $V^m(f) \coloneqq I_0^m \circ u : \mathcal{F} \to \mathbb{R}$. Again, by Lemma B.2, V^m is a well-defined functional over $\mathcal{F}$. Moreover, it represents $\succsim^m$. Indeed, for any $f, f' \in \mathcal{F}$:

$$\begin{aligned} f \succsim^m f' &\iff x_f \succsim^m x_{f'} \\ &\iff u(x_f) \geq u(x_{f'}) \\ &\iff I_0^m(u(f)) \geq I_0^m(u(f')) \\ &\iff V^m(f) \geq V^m(f') \; . \end{aligned}$$

Lemma B.5: *I_0^m is a normalized and concave niveloid.*

Proof: *Step 1: Monotonicity.* Take $\varphi, \psi \in B_0(\mathcal{G})$ and assume that $\varphi \geq \psi$. By Lemma B.2, we can find $f_\varphi, f_\psi \in \mathcal{F}$ such that $u(f_\varphi) = \varphi$ and $u(f_\psi) = \psi$. Then, for all $\omega' \in \Omega$,

$$u(f_\varphi(\omega')) = u(f_\varphi)(\omega') = \varphi(\omega') \geq \psi(\omega') = u(f_\psi)(\omega') = u(f_\psi(\omega'))$$

and, therefore, $f_\varphi(\omega) \succsim^m f_\psi(\omega)$. Then, since by Lemma B.4, $\succsim^m$ satisfies monotonicity and transitivity, $f_\varphi \succsim^m f_\psi$ and, therefore, $x_{f_\varphi} \succsim^m x_{f_\psi}$. We can, thus, conclude that

$$I_0^m(\varphi) = u(x_{f_\varphi}) \geq u(x_{f_\psi}) = I_0^m(\psi)$$

which proves the claim.

Step 2: Normalization. Take $k \in \mathbb{R}$. Since $\operatorname{Im} u = \mathbb{R}$, we can find $x^k \in X$ such that $u(x^k) = k$. Then:

$$I_0^m(k) = u(x^k) = k$$

showing that I_0^m is normalized.

Step 3: Translation Invariance. Take any $\varphi, \psi \in B_0(\mathcal{G})$ and $k, r \in \mathbb{R}$. By Lemma B.2 and surjectivity, we can find $f_\varphi, f_\psi \in \mathcal{F}$ and $x^k, x^r \in X$ such that $u(f_\varphi) = \varphi$, $u(f_\psi) = \psi$, $u(x^k) = k$, and $u(x^r) = r$. Now, for any $\alpha \in (0,1)$, since u is an affine operator, we have for each $\xi \in \{\varphi, \psi\}$, $l \in \{k, r\}$,

$$u(\alpha f_\xi + (1-\alpha)x^l) = \alpha u(f_\xi) + (1-\alpha)u(x^l) = \alpha\xi + (1-\alpha)l \ .$$

Moreover, by Lemma B.4 and the fact that $I_0^m \circ u$ represents $\succsim^m$:

$$\begin{aligned}
& I_0^m\left(\alpha\varphi + (1-\alpha)k\right) = I_0^m\left(\alpha\psi + (1-\alpha)k\right) \\
\Longrightarrow\ & I_0^m\left(u(\alpha f_\varphi + (1-\alpha)x^k)\right) = I_0^m\left(u(\alpha f_\psi + (1-\alpha)x^k)\right) \\
\Longrightarrow\ & \alpha f_\varphi + (1-\alpha)x^k \sim^m \alpha f_\psi + (1-\alpha)x^k \\
\Longrightarrow\ & \alpha f_\varphi + (1-\alpha)x^r \sim^m \alpha f_\psi + (1-\alpha)x^r \\
\Longrightarrow\ & I_0^m\left(u(\alpha f_\varphi + (1-\alpha)x^r)\right) = I_0^m\left(u(\alpha f_\psi + (1-\alpha)x^r)\right) \\
\Longrightarrow\ & I_0^m\left(\alpha\varphi + (1-\alpha)r\right) = I_0^m\left(\alpha\psi + (1-\alpha)r\right) \ .
\end{aligned}$$

Then, for any $\varphi', \psi' \in B_0(\mathcal{G})$ and $k', r' \in \mathbb{R}$, by letting $\varphi = \varphi'/\alpha$, $\psi = \psi'/\alpha$, $k = k'/(1-\alpha)$, and $r = r'/(1-\alpha)$ in the previous implication:

$$I_0^m(\varphi' + k') = I_0^m(\psi' + k') \implies I_0^m(\varphi' + r') = I_0^m(\psi' + r') \ .$$

Then, take any $\xi \in B_0(\mathcal{G})$ and $l \in \mathbb{R}$. By Step 2, I_0^m is normalized and, therefore, $I_0^m(\xi) = I_0^m(I_0^m(\xi))$. By what is shown above, this implies:

$$I_0^m(\xi + l) = I_0^m(I_0^m(\xi) + l) = I_0^m(\xi) + l$$

proving the claim.

Step 4: Quasi-concavity. Take any $\varphi, \psi \in B_0(\mathcal{G})$ such that $I_0^m(\varphi) = I_0^m(\psi)$ and $\alpha \in (0,1)$. By Lemma B.2, we can find $f_\varphi, f_\psi \in \mathcal{F}$ such that $\varphi = u(f_\varphi)$ and $\psi = u(f_\psi)$. Then:

$$V^m(f_\varphi) = I_0^m(u(f_\varphi)) = I_0^m(\varphi) = I_0^m(\psi) = I_0^m(u(f_\psi)) = V^m(f_\psi)$$

so that $f_\varphi \sim^m f_\psi$. Since $\succsim^m$ satisfies Axiom B.1, uncertainty aversion implies that

$$\alpha f_\varphi + (1-\alpha) f_\psi \succsim^m f_\psi$$

and, therefore:

$$\begin{aligned} I_0^m(\alpha\varphi + (1-\alpha)\psi) &= I_0^m(\alpha u(f_\varphi) + (1-\alpha)u(f_\psi)) \\ &= I_0^m(u(\alpha f_\varphi + (1-\alpha) f_\psi)) \\ &= V^m(\alpha f_\varphi + (1-\alpha) f_\psi) \\ &\geq V^m(f_\psi) \\ &= I_0^m(u(f_\psi)) = I_0^m(\psi) \end{aligned}$$

proving the claim.

By Steps 1-4 and Theorem 4 in Cerreia-Vioglio et al. (2014), it follows that I_0^m is a normalized and concave niveloid. $\square$

Denote by $I^m : B(\mathcal{G}) \to \mathbb{R}$ the unique normalized and concave niveloid extending I_0^m (see Lemma 25 in Maccheroni et al. (2006)). It is clear that $V^m = I^m \circ u$ on $\mathcal{F}$. Then, by Lemma 26 in Maccheroni et al. (2006), there exists a grounded, lower semicontinuous and convex function $c^m : \Delta \to [0,1]$ such that:

$$
\begin{aligned}
I^m(\varphi) &= \min_{p' \in \Delta(\mathcal{G})} \left\{ \int_\Omega \varphi dp' + c^m(p') \right\} \\
c^m(p) &= \sup_{\varphi' \in B(\mathcal{G})} \left\{ I^m(\varphi') - \int_\Omega \varphi' dp \right\}
\end{aligned} \tag{12}
$$

for all $\varphi \in B(\mathcal{G})$ and $p \in \Delta(\mathcal{G})$. Then, define $c(\cdot, m) \coloneqq c^m(\cdot)$ for all $m \in \mathcal{M}$. We have that for each $m \in \mathcal{M}$ and for each $f, f' \in \mathcal{F}$,

$$
\begin{aligned}
f \succsim^m f' &\iff V^m(f) \geq V^m(f') \\
&\iff I^m(u(f)) \geq I^m(u(f')) \\
&\iff \min_{p \in \Delta(\mathcal{G})} \left\{ \int_\Omega u(f) dp + c(p,m) \right\} \geq \min_{p \in \Delta(\mathcal{G})} \left\{ \int_\Omega u(f') dp + c(p,m) \right\},
\end{aligned}
$$

proving the representation in (5). We only need to check that $c(\cdot, m)$ is finite only on probabilities that are absolutely continuous with respect to m. This is the content of the next lemma.

Lemma B.6: *For all $m \in \mathcal{M}$, if $p \in \operatorname{dom} c(\cdot, m)$, then $p \ll m$ and $c(p,m) = 0$ if and only if $p = m$. In particular, c is a convex divergence.*

Proof of Lemma B.6: Fix any $m \in \mathcal{M}$.

We first show that if $p \in \operatorname{dom} c(\cdot, m)$, then p is absolutely continuous with respect to m. Suppose there exists a structured model $m \in \mathcal{M}$ and a $\hat{p} \in \operatorname{dom} c(\cdot, m)$ that is not absolutely continuous with respect to m. We show that $\succsim$ would violate Coherence. Indeed, we can find a measurable set $E \in \mathcal{G}$ such that $m(E) = 0$ but $\hat{p}(E) > 0$.

Consider the sequence of acts $(f_n)_{n\in\mathbb{N}} \subseteq \mathcal{F}$ such that for each $n \in \mathbb{N}$, $f_n = x_n E x_0$ where, since u is surjective, we can pick $x_0 \in u^{-1}(0)$ and $x_n \in u^{-1}(-n)$. Since $m(E) = 0$, $f_n = x_0$ $a.e.[m]$ for any $n \in \mathbb{N}$. Since $\hat{p} \in \operatorname{dom} c(\cdot, m)$, $c(\hat{p}, m) < \infty$, so that there exists $N \in \mathbb{N}$ large enough such that $c(\hat{p}, m) < N \cdot \hat{p}(E)$. Therefore,

$$\begin{aligned} V^m(f_N) = I^m(u(f_N)) &= \min_{p\in\Delta} \left\{ \int_\Omega u(f_N) dp + c(p, m) \right\} \\ &= \min_{p\in\Delta} \left\{ \int_E -N \; dp + c(p, m) \right\} \\ &= \min_{p\in\Delta} \{-N \; p(E) + c(p, m)\} \\ &\leq -N \; \hat{p}(E) + c(\hat{p}, m) \\ &< 0 = u(x_0) \end{aligned}$$

showing that $x_0 \succ^m f_N$ and, therefore, $x_0 E^m x_0 \not\sim f_N E^m x_0$. But since $x_0 = f_N$ with probability 1 according to m, this violates Coherence.

We now show that $c(p, m) = 0$ if and only if $p = m$. Let $P_0 \coloneqq \{p_0 \in \Delta(\Omega) : c(p_0, m) = 0\}$. First of all, P_0 is non-empty because $c(\cdot, m)$ is grounded. Moreover, $P_0 \subseteq \{p_0 \in \Delta(\Omega) : p_0 \ll m\}$ by what just shown above. Take $p_0 \ll m$ such that $p_0 \neq m$. Then, by Coherence there must exist $f \in \mathcal{F}$ such that $f E^m x \succsim x$, but $x \succ \int_\Omega f dp_0$. But, then,

$$\int_\Omega u(f) dp_0 + c(p_0, m) \geq \min_{p\in\Delta(\Omega)} \left\{ \int_\Omega u(f) + c(p, m) \right\} \geq u(x) > u\left(\int_\Omega u(f) dp_0 \right) = \int_\Omega u(f) dp_0$$

which implies that $c(p_0, m) > 0$. Since this holds for all $p_0 \ll m$ sich that $p_0 \neq m$, it must be the case that $\emptyset \neq P_0 \subseteq \{m\}$. That is, $c(p, m) = 0$ if and only if $p = m$.

□

As an almost immediate consequence of Lemma B.6, we show that for any $m \in \mathcal{M}$, if $\varphi, \psi \in B(\mathcal{G})$ and $\varphi = \psi$ a.e. $[m]$, then $I^m(\varphi) = I^m(\psi)$. Indeed:

$$m(\{\omega : \varphi(\omega) \neq \psi(\omega)\}) = 0 \implies \forall p \ll m,\ p(\{\omega : \varphi(\omega) \neq \psi(\omega)\}) = 0$$

and, therefore,

$$I^m(\varphi) = \min_{p \ll m} \left\{ \int_\Omega \varphi \ dp + c(p, m) \right\} = \min_{p \ll m} \left\{ \int_\Omega \psi \ dp + c(p, m) \right\} = I^m(\psi).$$

Finally, as far as uniqueness, that u is cardinally unique follows from Herstein and Milnor (1953). Moreover, the uniqueness of c given u is guaranteed by the fact that $\succsim^m$ is an unbounded variational preference and Proposition 6 in Maccheroni et al. (2006). ■

Nest, we show the characterization of the comparative notion of misspecificaiton aversion.

Proof of Proposition 2: Suppose that $\succsim_1$ and $\succsim_2$ are two misspecification averse preferences. Let (u_1, c_1) and (u_2, c_2) represent respectively $(\succsim_2^m)_{m \in \mathcal{M}}$ and $(\succsim_2^m)_{m \in \mathcal{M}}$ as in Proposition 1 and define I_1^m and I_2^m accordingly for all $m \in \mathcal{M}$. Suppose that u_2 is a positive affine transformation of u_1 and $c_1 \leq c_2$. Without loss of generality, assume that $u_1 = u_2 = u$. Fix any $m \in \mathcal{M}$ and take any $f \in \mathcal{F}$ and $x \in X$ such that $fE^m x \succsim_1^m x$. Then, $f \succsim_1^m x$ and, therefore, $I_1^m(u(f)) \geq u(x)$. Then:

$$\begin{aligned} I_2^m(u(f)) &= \min_{p \in \Delta} \left\{ \int_\Omega u(f) dp + c_2(p, m) \right\} \\ &\geq \min_{p \in \Delta} \left\{ \int_\Omega u(f) dp + c_1(p, m) \right\} \\ &= u(x) \end{aligned}$$

so that $f \succsim_2^m x$, and, therefore, $fE^m x \succsim_2 x$.

As for the other direction, note that Equation 7 and nontriviality imply that u_2 is a positive affine transformation of u_1. Without loss of generality, set $u_1 = u_2 = u$. Fix any $m \in \mathcal{M}$ and take $\varphi \in B_0(\mathcal{G})$. Let $f \in \mathcal{F}$ be such that $u(f) = \varphi$ and $x \in X$ such that $f \sim_1^m x$. Then, condition 7 implies that $f \succsim_2^m x$, so that

$$I_1^m(\varphi) = I_1^m(u(f)) = u(x) \leq I_2^m(u(f)) = I_2^m(\varphi).$$

Therefore, $I_1^m(\varphi) \leq I_2^m(\varphi)$ for all $\varphi \in B_0(\mathcal{G})$. Since the latter is dense in the space $B(\mathcal{G})$, we conclude that $I_1 \leq I_2$. Then, using Equation (12):

$$\begin{aligned} c_1(p,m) &= \sup_{\varphi' \in B(\mathcal{G})} \left\{ I_1^m(\varphi') - \int_\Omega \varphi' dp \right\} \\ &\leq \sup_{\varphi' \in B(\mathcal{G})} \left\{ I_2^m(\varphi') - \int_\Omega \varphi' dp \right\} = c_2(p,m) \end{aligned}$$

for all $p \in \Delta$. ■

We conclude this section by proving the existence of a generalized conditional expectation.

Proof of Corollary 3: First, we show that for any given $\varphi \in B(\mathcal{G})$, $I^m(\varphi)$ is measurable as a function of m.

Lemma B.7: *The map $m \mapsto I^m(\varphi)$ is a $\mathcal{D}_{\mathcal{M}}$-measurable and bounded functional for all $\varphi \in B(\Omega, \mathcal{G})$.*

Proof of Lemma B.7: Fix $\varphi \in B(\Omega, \mathcal{G})$ arbitrarily. We first show that $m \mapsto I^m(\varphi)$ is bounded. Indeed, since φ is bounded, there exist $k, K \in \mathbb{R}$ such that $k \leq \varphi \leq K$. By Lemma B.5, for each $m \in \mathcal{M}$, I^m is normalized and monotone and, therefore,

$$k = I^m(k) \leq I^m(\varphi) \leq I^m(K) = K$$

proving boundedness. We now show that $m \mapsto I^m(\varphi)$ is also measurable. Take any real number $r \in \mathbb{R}$. We want to show that $\{m \in \mathcal{M} : I^m(\varphi) > r\}$ is a measurable set in $\mathcal{D}_{\mathcal{M}}$. Since u is surjective, take x_r such that $u(x_r) = r$. Moreover, by Lemma B.2, we can pick f_φ such that $u(f_\varphi) = \varphi$. Then, we have:

$$\{m \in \mathcal{M} : I^m(\varphi) > r\} = \{m \in \mathcal{M} : I^m(u(f_\varphi)) > u(x_r)\}$$
$$= \{m \in \mathcal{M} : f_\varphi E^m x_r \succsim x_r\}$$

and the latter is measurable since $\succsim$ satisfies Coherence. This proves that $m \mapsto I^m(\varphi)$ is bounded and measurable for any $\varphi \in B(\Omega, \mathcal{G})$. ■

Denote by $\mathfrak{q}_0$ the restriction of $\mathfrak{q}$ to Ω_0. Clearly, $\mathfrak{q}_0$ is $\mathcal{A}_{\Omega_0}/\mathcal{D}_{\mathcal{M}}$, where $\mathcal{A}_{\Omega_0}$ is the relative σ-algebra $\mathcal{A} \cap \Omega_0$. Fix any $\varphi \in B(\Omega, \mathcal{G})$. Since $m \mapsto I^m(\varphi)$ is bounded and $\mathcal{D}_{\mathcal{M}}$-measurable by Lemma B.7, it follows that the composition

$$\begin{array}{rccccc} I^{\mathfrak{q}(\cdot)}(\varphi): & (\Omega_0, \mathcal{A}_{\Omega_0}) & \to & (\mathcal{M}, \mathcal{D}_{\mathcal{M}}) & \to & (\mathbb{R}, \mathcal{B}(\mathbb{R})) \\ & \omega & \mapsto & \mathfrak{q}(\omega) & \mapsto & I^{\mathfrak{q}(\omega)}(\varphi) \end{array}$$

is a $\mathcal{A}_{\Omega_0}$-measurable and bounded functional. Obtain $I_{\mathcal{A}}(\varphi)$ by extending $I^{\mathfrak{q}(\cdot)}(\varphi)$ to the whole Ω in the following way: $I_{\mathcal{A}}(\varphi)(\omega) = I^{\mathfrak{q}(\omega)}\varphi$ if $\omega \in \Omega_0$ and $I_{\mathcal{A}}(\varphi)(\omega) = 0$ if $\omega \in \Omega \setminus \Omega_0$. It is easy to see that $I_{\mathcal{A}}(\varphi) \in B(\mathcal{A})$. Moreover, take any $A \in \mathcal{A}$ and fix $m \in \mathcal{M}$ arbitrarily. We know that $m(E^m) = 1$, so that $m(\Omega \setminus E^m) = 0$, where we recall that $E^m = \{\omega \in \Omega : \mathfrak{q}(\omega) = m\}$. By Lemma B.6, we also have that if $p \in \operatorname{dom} c(\cdot, m)$, it must be the case that p is absolutely continuous with respect to m. Then, $p(\Omega \setminus E^m) = 0$ and $p(E^m) = 1$ for all $p \in \operatorname{dom} c(\cdot, m)$. Moreover, since $A \in \Lambda$ by Lemma B.1, we have that either $m(A) = 1$ or $m(A) = 0$. In any case, this implies that for any $p \in \operatorname{dom} c(\cdot, m)$,

$$p(A \cap E^m) = p(A)p(E^m) = p(A) = m(A).$$

Then:

$$\begin{aligned}
I^m(I_{\mathcal{A}}(\varphi)\chi_A) &= \min_{p\in\Delta}\left\{\int_\Omega I_{\mathcal{A}}(\varphi)(\omega)\chi_A(\omega)\ dp(\omega)+c(p,m)\right\}\\
&= \min_{p\in \operatorname{dom} c(\cdot,m)}\left\{\int_{A\cap E^m} I^{\mathfrak{q}(\omega)}(\varphi)\ dp(\omega)+c(p,m)\right\}\\
&= \min_{p\in \operatorname{dom} c(\cdot,m)}\left\{\int_{A\cap E^m} I^{m}(\varphi)\ dp(\omega)+c(p,m)\right\}\\
&= \min_{p\in \operatorname{dom} c(\cdot,m)}\left\{I^q(\varphi)\ q(A)+c(p,m)\right\}\\
&= I^m(\varphi)\ m(A)\\
&= I^m(\varphi\chi_A).
\end{aligned}$$

The last equality follows from the fact that $m(A)\in\{0,1\}$. Indeed, if $m(A)=0$,

$$I^m(\varphi\chi_A)=\min_{p\in \operatorname{dom} c(\cdot,m)}\left\{\int_A \varphi dp+c(p,m)\right\}=0=I^m(\varphi)m(A)$$

and if $m(A)=1$,

$$\begin{aligned}
I^m(\varphi\chi_A) &= \min_{p\in \operatorname{dom} c(\cdot,m)}\left\{\int_A \varphi dp+c(p,m)\right\}\\
&= \min_{p\in \operatorname{dom} c(\cdot,m)}\left\{\int_\Omega \varphi dp+c(p,m)\right\}=I^m(\varphi)m(A)\ .
\end{aligned}$$

■

Appendix D. Structured Functionals

Throughout the section, assume that $(\Omega,\mathcal{G},\mathcal{M})$ is a structured space with sufficient sub σ-algebra $\mathcal{A}\subseteq\mathcal{G}$ and best-fit map $\mathfrak{q}:\Omega\to\Delta^\sigma(\mathcal{G})$ satisfying the properties in Definition 1, that I^m is given as in the representation of Proposition 1 for all structured models $m\in\mathcal{M}$, and that $I_{\mathcal{A}}$ is the common generalized conditional expectation of $\mathcal{M}$ given $\mathcal{A}$, which exists by Corollary 3. Notice that for each $\varphi\in B(\mathcal{G}$, we can see

$I^m(\varphi)$ as a function from structured models to $\mathbb{R}$:

$$I(\varphi,\cdot):\mathcal{M}\to\mathbb{R},\quad m\mapsto I(\varphi,m):=I^m(\varphi).$$

Define the operator $T:B(\Omega,\mathcal{G})\to\mathbb{R}^{\mathcal{M}}$ such that for all $\varphi\in B(\Omega,\mathcal{G})$,

$$T(\varphi)(m)=I(\varphi,m)$$

for all $m\in\mathcal{M}$. By Lemma B.7, we have that $\operatorname{Im}T\subseteq B(\mathcal{M},\mathcal{D}_{\mathcal{M}})$.

Lemma B.8: *Let $T(B(\mathcal{A}))$ and $T(B_0(\mathcal{A}))$ be the images through T of $B(\mathcal{A})$ and $B_0(\mathcal{A})$ respectively. Then, $T(B(\mathcal{A}))=\operatorname{Im}T$ and $T(B_0(\mathcal{A}))$ is supnorm dense in $\operatorname{Im}T$. Moreover, T preserves lattice operations when restricted to $B(\mathcal{A})$. In particular, $\operatorname{Im}T$ is a lattice.*

Proof of Lemma B.8: It is clear that

$$\begin{aligned}T(B(\mathcal{A}))&=\{I(\varphi,\cdot):\varphi\in B(\Omega,\mathcal{A})\}\\&\subseteq\{I(\varphi,\cdot):\varphi\in B(\Omega,\mathcal{G})\}=\operatorname{Im}T\end{aligned}$$

since $\mathcal{A}$ is a a sub-σ-algebra of $\mathcal{G}$. As for the reverse inclusions, take any $\xi\in\operatorname{Im}T$ and let $\varphi_\xi\in B(\mathcal{G})$ be such that $\xi=T(\varphi_\xi)$. Then, by Corollary 3, $I_{\mathcal{A}}(\varphi_\xi)\in B(\mathcal{A})$ and for all $m\in\mathcal{M}$,

$$T(I_{\mathcal{A}}(\varphi_\xi))(m)=I^m\left(I_{\mathcal{A}}(\varphi_\xi)\right)=I^m(\varphi_\xi)=T(\varphi_\xi)(m)=\xi(m),$$

so that $\xi\in T(B(\mathcal{A}))$, showing that $\operatorname{Im}T\subseteq T(B(\mathcal{A}))$ Next, we show that $T(B_0(\mathcal{A}))$ is supnorm dense in $\operatorname{Im}T$. Take $\xi\in\operatorname{Im}T$ and a corresponding $\varphi_\xi\in B(\mathcal{A})$ such that $\xi=T(\varphi_\xi)$ (which exists given what shown above). Since $B_0(\mathcal{A})$ is supnorm dense in $B(\mathcal{A})$, we can find a sequence $(\varphi_n)_n\subseteq B_0(\mathcal{A})$ such that $||\varphi_n-\varphi_\xi||_\infty\to 0$. Define $\xi_n=T(\varphi_n)$ for each $n\in\mathbb{N}$ and note that $(\xi_n)_n\subseteq T(B_0(\mathcal{A}))$. We show that ξ_n

converges to ξ in the supnorm. Indeed, for each $m \in \mathcal{M}$, since I^m is a niveloid and, therefore, Lipschitz continuous, we have that:

$$|\xi(m) - \xi_n(m)| = |T(\varphi)(m) - T(\varphi_n)(m)| = |I^m(\varphi) - I^m(\varphi_n)| \leq ||\varphi - \varphi_n||_\infty$$

and, therefore,

$$||\xi - \xi_n||_\infty = \sup_{m \in \mathcal{M}} |\xi(m) - \xi_n(m)| \leq ||\varphi - \varphi_n||_\infty \to 0.$$

Finally, we show that T preserves lattice operations on $B(\mathcal{A})$. Indeed, pick $\varphi, \tilde{\varphi} \in B(\mathcal{A})$ arbitrarily. Since $B_0(\mathcal{A})$ is supnorm dense in $B(\mathcal{A})$, we can take sequences $(\varphi)_n, (\tilde{\varphi}_n)_n \subseteq B_0(\mathcal{A})$ such that $||\varphi - \varphi_n||_\infty, ||\tilde{\varphi} - \tilde{\varphi}_n||_\infty \to 0$. For each $n \in \mathbb{N}$, we can find a finite partition $(E_n^i)_{i=1}^k$ and reals $(r_n^i)_{i=1}^k$, $(\tilde{r}_n^i)_{i=1}^k$ such that:

$$\varphi_n = \sum_{i=1}^k \chi_{E_n^i} r_n^i, \quad \tilde{\varphi}_n = \sum_{i=1}^k \chi_{E_n^i} \tilde{r}_n^i.$$

Fix any $m \in \mathcal{M}$. By Lemma B.1, for each $n \in \mathbb{N}$, there is a unique E_n^l in the partition such that $m(E_n^l) = 1$. Therefore, $\varphi_n = r_n^l$ and $\tilde{\varphi}_n = \tilde{r}_n^l$ a.e. $[m]$, so that by Proposition 1 and normalization, $I^m(\varphi_n) = I^m(r_n^l) = r_n^l$ and $I^m(\tilde{\varphi}_n) = I^m(\tilde{r}_n^l) = \tilde{r}_n^l$ for all $n \in \mathbb{N}$. Clearly, it is also the case that $\varphi_n \vee \tilde{\varphi}_n = r_n^l \vee \tilde{r}_n^l$ a.e. $[m]$ so that $I^m(\varphi_n \vee \tilde{\varphi}_n) = I^m(r_n^l \vee \tilde{r}_n^l) = r_n^l \vee \tilde{r}_n^l$ for all $n \in \mathbb{N}$. Therefore:

$$I^m(\varphi_n \vee \tilde{\varphi}_n) = r_n^l \vee \tilde{r}_n^l = I^m(\varphi_n) \vee I^m(\tilde{\varphi}_n)$$

for all $n \in \mathbb{N}$. Since lattice operations are continuous and I^m is Lipschitz, taking limits, it follows that

$$T(\varphi \vee \tilde{\varphi})(m) = I^m(\varphi \vee \tilde{\varphi}) = I^m(\varphi) \vee I^m(\tilde{\varphi}) = T(\varphi)(m) \vee T(\tilde{\varphi})(m)$$

Since m was chosen arbitrarily, we can conclude that $T(\varphi \vee \tilde{\varphi}) = T(\varphi) \vee T(\tilde{\varphi})$. That $\operatorname{Im} T$ is a lattice follows from the fact that $\operatorname{Im} T = T(B(\mathcal{A}))$ and $T|_{B(\mathcal{A})}$ preserves lattice operations. ■

Recall that $B_0(\mathcal{D}_\mathcal{M}) \coloneqq B_0(\mathcal{M}, \mathcal{D}_\mathcal{M})$ and $B(\mathcal{D}_\mathcal{M}) \coloneqq B(\mathcal{M}, \mathcal{D}_\mathcal{M})$ are, respectively, the spaces of simple and bounded functions on the set of structured models $\mathcal{M}$ measurable with respect to $\mathcal{D}_\mathcal{M}$. The following result shows that these spaces can be covered by applying the operator T respectively to $B_0(\mathcal{A})$ and $B(\mathcal{A})$. Further, characteristic functions of sets in $\mathcal{D}_\mathcal{M}$ can be recovered by applying the operator T to characteristic functions of sets in $\mathcal{A}$.

Lemma B.9: $\operatorname{Im} T = B(\mathcal{M}, \mathcal{D}_\mathcal{M})$. *Moreover,* $T(B_0(\mathcal{A})) = B_0(\mathcal{M}, \mathcal{D}_\mathcal{M})$ *and* $T(\{\chi_E : E \in \mathcal{A}\}) = \{\chi_D : D \in \mathcal{D}_\mathcal{M}\}$.

Proof of Lemma B.9: We prove the results via a series of steps.

Step (i). For all $E \in \mathcal{A}$*, there exists* $D_E \in \mathcal{D}_\mathcal{M}$ *such that* $T(\chi_E) = \chi_{D_E}$.

Proof: Take any $E \in \mathcal{A}$. By Lemma B.1, $E \in \Lambda$ and, thererfore, for all $m \in \mathcal{M}$, either $m(E) = 1$ or $m(E) = 0$. But then for all $m \in \mathcal{M}$:

$$m(E) = 1 \implies \chi_E = 1 \text{ a.e. } [m] \implies T(\chi_E)(m) = I^m(\chi_E) = I^m(1) = 1,$$

$$m(E) = 0 \implies \chi_E = 0 \text{ a.e. } [m] \implies T(\chi_E)(m) = I^m(\chi_E) = I^m(0) = 0.$$

Therefore, $\operatorname{Im} T(\chi_E) \in \{0, 1\}$. Moreover, by Lemma B.7, $D_E \coloneqq [T(\chi_E)]^{-1}(\{1\}) \in \mathcal{D}_\mathcal{M}$ and $T(\chi_E) = \chi_{D_E}$ as we wanted to show. □

Step (ii). For all $D \in \mathcal{D}_\mathcal{M}$*, there exists* $E^D \in \mathcal{A}$ *such that* $T(\chi_{E^D}) = \chi_D$.

Proof: Take any $D \in \mathcal{D}_\mathcal{M}$ and let $E^D = \mathfrak{q}^{-1}(D)$. Since the space is structured, $E^D \in \mathcal{A}$ and $m(E^D) = 1$ if $m \in D$ and $m(E^D) = 0$ if $m \in \mathcal{M} \setminus D$. But then for all

$m \in \mathcal{M}$:

$$m \in D \implies \chi_{E^D} = 1 \text{ a.e. } [m] \implies T(\chi_{E^D})(m) = I^m(\chi_{E^D}) = I^m(1) = 1,$$

$$m(E) \in \mathcal{M} \setminus D \implies \chi_{E^D} = 0 \text{ a.e. } [m] \implies T(\chi_{E^D})(m) = I^m(\chi_{E^D}) = I^m(0) = 0,$$

and we can, thus, conclude that $T(\chi_{E^D}) = \chi_D$. □

Steps (i) and (ii) together imply that $T(\{\chi_E : E \in \mathcal{A}\}) = \{\chi_D : D \in \mathcal{D}_{\mathcal{M}}\}$.

Step (iii). $T(B_0(\mathcal{A})) \subseteq B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$.

Proof: Take $\xi \in T(B_0(\mathcal{A}))$. By definition, there exists $\varphi_\xi \in B_0(\mathcal{A})$ such that $\xi = T(\varphi_\xi)$. Then, there exists a partition $(E_i)_{i=1}^k \subseteq \mathcal{A}$ and reals $(r_i)_{i=1}^k$ such that $\varphi_\xi = \sum_{i=1}^k \chi_{E_i} r_i$. By Step (i), we have that for each $i = 1, \ldots, k$, we can find $D_{E_i} \in \mathcal{D}_{\mathcal{M}}$ such that $T(\chi_{E_i}) = \chi_{D_{E_i}}$. Moreover, since for all $i = 1, \ldots, k$, $E_i \in \mathcal{A} \subseteq \Lambda$ by Lemma B.1, either $m(E_i) = 1$ or $m(E_i) = 0$ for each $m \in \mathcal{M}$. It follows that for each m, there is a unique element in the partition E_{j_m} such that $m(E_{j_m}) = 1$ and $m(E_i) = 0$ if $i \neq j_m$. Then, for each $m \in \mathcal{M}$,

$$\varphi_\xi = r_{j_m} \text{ a.e. } [m] \implies T(\varphi_\xi)(m) = I^m(\varphi_\xi) = I^m(r_{j_m}) = r_{j_m}$$

and, since $\chi_{E_{j_m}} = 1$ a.e. $[m]$ and $\chi_{E_i} = 0$ a.e. $[m]$ for $i \neq j_m$,

$$\chi_{D_{E_{j_m}}}(m) = T(\chi_{E_{j_m}})(m) = I^m(\chi_{E_{j_m}}) = I^m(1) = 1 \implies m \in D_{E_{j_m}}$$

$$\forall i \neq j_m, \quad \chi_{D_{E_i}}(m) = T(\chi_{E_i})(m) = I^m(\chi_{E_i}) = I^m(0) = 0 \implies m \notin D_{E_i}.$$

It follows that $\varphi_\xi = \sum_{i=1}^k \chi_{D_{E_i}} r_i \in B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$. □

Step (iv). $B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \subseteq T(B_0(\mathcal{A}))$, In particular, for all $D \in \mathcal{D}_{\mathcal{M}}$, there exists $E^D \in \mathcal{A}$ such that $\chi_D = T(\chi_{E^D})$.

Proof: Take any $\xi \in B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$. By definition, there exists a partition $(D_i)_{i=1}^k \subseteq \mathcal{D}_{\mathcal{M}}$ of $\mathcal{M}$ and reals $(r_i)_{i=1}^k$ such that $\xi = \sum_{i=1}^k \chi_{D_i} r_i$. By Step (ii), for each $i =$

$1, \dots, k$, we can find $E^{D_i} \in \mathcal{A}$ such that $\chi_{D_i} = T(\chi_{E^{D_i}})$. Define $\varphi_\xi := \sum_{i=1}^k \chi_{E^{D_i}} r_i$. Clearly, $\varphi_\xi \in B_0(\mathcal{A})$. Moreover, for each $m \in \mathcal{M}$, let D_{j_m} be the unique element of the partition such that $m \in D_{j_m}$. We know by Lemma B.1 that since $E^{D_{j_m}} \in \mathcal{A}$, $m(E^{D_{j_m}}) \in \{0, 1\}$. If $m(E^{D_{j_m}}) = 0$, then $\chi_{E^{D_{j_m}}} = 0$ a.e. $[m]$ and, therefore, $T(E^{D_{j_m}})(m) = I^m(E^{D_{j_m}}) = I^m(0) = 0 \neq \chi_{D_{j_m}}(m) = 1$, a contradiction. We conclude that $m(E^{D_{j_m}}) = 1$ so that $\varphi_\xi = r_{j_m}$ a.e. $[m]$. Therefore,

$$T(\varphi_\xi)(m) = I^m(\varphi_\xi) = I^m(r_{j_m}) = r_{j_m} = r_{j_m}\chi_{D_{j_m}}(m) = \xi(m).$$

for all $m \in \mathcal{M}$. It follows that $T(\varphi_\xi) = \xi$, showing that $B_0(\mathcal{M}, \mathcal{D}_\mathcal{M}) \subseteq T(B_0(\mathcal{A}))$. □

Step (iii) and (iv) imply that $B_0(\mathcal{M}, \mathcal{D}_\mathcal{M}) = T(B_0(\mathcal{A}))$. Then, we have the following chain of inclusions:

$$B_0(\mathcal{M}, \mathcal{D}_\mathcal{M}) \subseteq T(B_0(\mathcal{A})) \subseteq B(\mathcal{M}, \mathcal{D}_\mathcal{M}).$$

Moreover, $B_0(\mathcal{M}, \mathcal{D}_\mathcal{M})$ is supnorm dense in $B(\mathcal{M}, \mathcal{D}_\mathcal{M})$ and by Lemma B.8, $T(B_0(\mathcal{A}))$ is supnorm dense in $\operatorname{Im} T$. Taking the supnorm closure of the previous chain of inclusions, we obtain that:

$$B(\mathcal{M}, \mathcal{D}_\mathcal{M}) = \operatorname{cl} B_0(\mathcal{M}, \mathcal{D}_\mathcal{M}) \subseteq \operatorname{cl} T(B_0(\mathcal{A})) = \operatorname{Im} T \subseteq \operatorname{cl} B(\mathcal{M}, \mathcal{D}_\mathcal{M}) = B(\mathcal{M}, \mathcal{D}_\mathcal{M})$$

and, therefore, we can conclude that $\operatorname{Im} T = B(\mathcal{M}, \mathcal{D}_\mathcal{M})$. ■

Lemma B.10:

(i) *If $\xi, \xi' \in B_0(\mathcal{M}, \mathcal{D}_\mathcal{M})$ are such that $\xi \geq \xi'$, then there exist $\varphi_\xi, \varphi_{\xi'} \in B_0(\mathcal{A})$ such that $\varphi_\xi \geq \varphi_{\xi'}$ and $\xi = T(\varphi_\xi)$, $\xi' = T(\varphi_{\xi'})$.*

(ii) *If $(\xi_n)_n \subseteq T(B_0(\mathcal{A}))$ is an increasing (decreasing) sequence uniformly bounded above (below) by a constant K, there exists an increasing (decreasing) sequence $(\varphi_n)_n \subseteq B_0(\mathcal{A})$ such that $\xi_n = T(\varphi_n)$ and $\varphi_n \leq K$ $(\varphi_n \geq K)$ for all $n \in \mathbb{N}$.*

(iii) If $\xi \in \operatorname{Im} T$ and $(\xi_n)_n \subseteq T(B_0(\mathcal{A}))$ such that $\xi_n \uparrow \xi$ ($\xi_n \downarrow \xi$), then we can find an increasing (decreasing) sequence $(\varphi_n)_n \subseteq B_0(\mathcal{A})$ and $\varphi \in B(\mathcal{A})$ such that $\varphi_n \uparrow \varphi$ ($\varphi_n \downarrow \varphi$), $\xi = T(\varphi)$, and $\xi_n = T(\varphi_n)$ for all $n \in \mathbb{N}$. Moreover, if $K \in \mathbb{R}$ and $\xi \leq K$ ($\xi \geq K$), then $\varphi \leq K$ ($\varphi \geq K$).

PROOF OF LEMMA B.10: We prove the lemma in a number of steps.

PROOF OF (i): Take $\xi, \xi' \in T(B_0(\mathcal{A}))$ such that $\xi \geq \xi'$. By definition, we can pick $\varphi_\xi, \varphi_{\xi'} \in B_0(\mathcal{A})$ such that $\xi = T(\varphi_\xi)$ and $\xi' = T(\varphi_{\xi'})$. Moreover, we can find a partition $(E_i)_{i=1}^n \subseteq \mathcal{A}$ of Ω and reals $(r_i)_{i=1}^n, (r_i')_{i=1}^n$ such that

$$\varphi_\xi = \sum_{i=1}^n \chi_{E_i} r_i, \quad \varphi_{\xi'} = \sum_{i=1}^n \chi_{E_i} r_i'.$$

Take an element E_k in the partition. If $m(E_k) = 0$ for all $m \in \mathcal{M}$, we can assume wlog that $r_k = r_k'$. Indeed, for all $m \in \mathcal{M}$, $\varphi_{\xi'} = \sum_{i \neq k} \chi_{E_i} r_i' + \chi_{E_k} r_k$ a.e. $[m]$ and by Proposition 1, this implies

$$\xi' = T(\varphi_{\xi'})(m) = I^m(\varphi_{\xi'}) = I^m(\sum_{i \neq k} \chi_{E_i} r_i' + \chi_{E_k} r_k) = T(\sum_{i \neq k} \chi_{E_i} r_i' + \chi_{E_k} r_k)(m).$$

If there exists $m \in \mathcal{M}$ such that $m(E_k) \neq 0$, then $m(E_k) = 1$ since $E_k \in \mathcal{A} \subseteq \Lambda$ by Lemma B.1. Therefore, $\varphi_\xi = r_k$ and $\varphi_{\xi'} = r_k'$ a.e. $[m]$ and, therefore:

$$r_k = I^m(r_k) = I^m(\varphi_\xi) = T(\varphi_\xi)(m) = \xi(m),$$
$$r_k' = I^m(r_k') = I^m(\varphi_{\xi'}) = T(\varphi_{\xi'})(m) = \xi'(m),$$

and, we conclude that $r_k = \xi(m) \geq \xi'(m) = r_k'$. We have thus shown that $r_i \geq r_i'$ for all $i = 1, \ldots, n$. Hence, it follows that $\varphi_\xi \geq \varphi_{\xi'}$. It is then immediate to see that since each I^m is normalized, if $\xi \leq K$ for some K in $\mathbb{R}$, we can find $\varphi_\xi \in B_0(\mathcal{A})$ such that $\varphi_\xi \leq K$ and $\xi = T(\varphi_\xi)$. □

Proof of (ii): Take a sequence $(\xi_n)_n \subseteq T(B_0(\mathcal{A}))$ and $K \in \mathbb{R}$ such that $\xi_n \leq \xi_{n+1} \leq K$ for all $n \in \mathbb{N}$. By Step (i), we can find a sequence $\varphi_{\xi_n} \in B_0(\mathcal{A})$ such that $\xi_n = T(\varphi_{\xi_n})$ and $\varphi_{\xi_n} \leq K$ for all $n \in \mathbb{N}$. However, this sequence is not necessarily increasing. Then, define for each $n \in \mathbb{N}$, $\varphi_n(\omega) = \sup_{k \leq n} \varphi_{\xi_k}(\omega)$ for all ω. Notice that $\varphi_n : \Omega \to \mathbb{R}$ is well-defined and in $B_0(\mathcal{A})$. Moreover, the sequence $(\varphi_n)_n$ so constructed is increasing and uniformly bounded above by K. Moreover, since T preserves lattice operations by Lemma B.8, we have that for each $n \in \mathbb{N}$,

$$T(\varphi_n) = T\left(\sup_{k \leq n} \varphi_{\xi_k}\right) = \sup_{k \leq n} T\left(\varphi_{\xi_k}\right) = \sup_{k \leq n} \xi_k = \xi_n,$$

where the last equality follows from the fact that $(\xi_n)_n$ is a monotonically increasing sequence. □

Proof of (iii): Take a sequence $(\xi_n)_n \subseteq T(B_0(\mathcal{A}))$ and $\xi \in \operatorname{Im} T$ such that $\xi_n \uparrow \xi$. Since ξ is bounded, $K_0 = \sup_{m \in \Omega} \xi$ is finite. Moreover, we have that $\xi_n \leq \xi \leq K_0$ for all $n \in \mathbb{N}$. By by point (ii), we can find an increasing sequence $(\varphi_n)_n \subseteq B_0(\mathcal{A})$ such that $\xi_n = T(\varphi_n)$ and $\varphi_n \leq K_0$ for all $n \in \mathbb{N}$. Since for each $\omega \in \Omega$, $(\varphi_n(\omega))_n$ is a monotonically increasing sequence of numbers bounded above by K_0, it converges to some $\lim_n \varphi_n(\omega) \leq K_0$. Therefore, the pointwise limit $\varphi := \lim_n \varphi_n$ is well-defined, it is in $B(\mathcal{A})$, and it is uniformly bounded above by K_0. Moreover, we have that for all $n \in \mathbb{N}$,

$$k = \min_{\omega \in \Omega} \varphi_1(\omega) \leq \varphi_1 \leq \varphi_n \leq K_0 \implies ||\varphi_n||_\infty \leq \max\{|k|, |K_0|\}.$$

Therefore, $(\varphi_n)_n$ is uniformly bounded in the norm. Moreover, for each $m \in \mathcal{M}$, Thereom 13 in Maccheroni et al. (2006) and Proposition 5 in Cerreia-Vioglio et al. (2014), imply that I^m has the Lebesgue property. Therefore:

$$T(\varphi)(m) = I^m(\varphi) = I^m(\lim_n \varphi_n) = \lim_n I^m(\varphi_n) = \lim_n \xi_n(m) = \xi(m).$$

It is immediate to see that for all $k \in \mathbb{R}$ such that $\xi \leq K$, $K \geq K_0$ and, therefore, $\varphi \leq K$. □

This concludes the proof of the lemma. ■

PROPOSITION B.11: *The following are equivalent:*

(i) $I : B(\mathcal{A}) \to \mathbb{R}$ *is normalized, monotone, and such that for all* $\varphi, \varphi' \in B_0(\mathcal{A})$,

$$(\forall m \in \mathcal{M},\ I^m(\varphi) \geq I^m(\psi)) \implies I(\varphi) \geq I(\psi).$$

(ii) there exists a normalized and monotone functional $\hat{I} : B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ *such that for all* $\varphi \in B_0(\mathcal{A})$,

$$I(\varphi) = \hat{I}(T(\varphi)).$$

Moreover, $\hat{I}$ *is unique and*

- $\hat{I}$ *is continuous if and only if* I *is continuous.*
- $\hat{I}$ *is quasiconcave if and only if* I *is quasiconcave.*
- $\hat{I}$ *is monotone continuous if and only if* I *is monotone continuous.*

PROOF OF PROPOSITION B.11:

(i) *implies* (ii). Define $\hat{I} : B_0(\mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ as follows: for all $\xi \in B_0(\mathcal{D}_{\mathcal{M}})$,

$$\hat{I}(\xi) = I(\varphi_\xi),$$

where $\varphi_\xi \in B_0(\mathcal{A})$ is chosen so that $\xi = T(\varphi_\xi)$.

Step 1: $\hat{I}$ *is well-defined.* Pick $\xi \in B_0(\mathcal{D}_{\mathcal{M}})$ arbitrarily. That a $\varphi_\xi \in B_0(\mathcal{A})$ such that $\xi = T(\varphi_\xi)$ exists follows from Lemma B.9. Moreover, suppose there are two $\varphi, \psi \in B_0(\mathcal{A})$ such that $T(\varphi)(m) = I^m(\varphi) = \xi(m) = I^m(\psi) = T(\psi)(m)$ for all $m \in \mathcal{M}$. Then, by assumption, it must be the case that $I(\varphi) = I(\psi)$, showing that $\hat{I}$ is well-defined.

Step 2: $\hat{I}$ is normalized. Take any $k \in \mathbb{R}$. Then, since each I^m is normalized, it follows that $k = I^m(k) = T(k)(m)$ for all $m \in \mathcal{M}$. By definition, it follows that $\hat{I}(k) = I(k) = k$, where the last equality follows from the assumption that I is normalized. This proves the step.

Step 3: $\hat{I}$ is monotone. Take $\xi, \xi' \in \operatorname{Im} T$ such that $\xi \geq \xi'$. By Lemma B.9, $\xi, \xi' \in T(B_0(\mathcal{A}))$ and, therefore, Lemma B.10 implies that we can find $\varphi_\xi, \varphi_{\xi'} \in B_0(\mathcal{A})$ such that $\varphi_\xi \geq \varphi_{\xi'}$ and $\xi = T(\varphi_\xi)$, $\xi' = T(\varphi_{\xi'})$. Since I is monotone

$$\hat{I}(\xi) = \hat{I}(T(\varphi_\xi)) = I(\varphi_\xi) \geq I(\varphi_{\xi'}) = \hat{I}(T(\varphi_{\xi'})) = \hat{I}(\xi')$$

showing that also $\hat{I}$ is monotone.

Step 4: $\hat{I}$ is unique. Suppose there is another $\tilde{I} : B_0(\mathcal{M}, \mathcal{D}_\mathcal{M}) \to \mathbb{R}$ such that $I(\varphi) = \tilde{I}(T(\varphi))$ for all $\varphi \in B_0(\mathcal{A})$. Then, take any $\xi \in B_0(\mathcal{M}, \mathcal{D}_\mathcal{M})$. By Lemma B.9, there exists $\varphi_\xi \in B_0(\mathcal{A})$ and such that $\xi = T(\varphi_\xi)$. Then,

$$\tilde{I}(\xi) = \tilde{I}(T(\varphi_\xi)) = I(\varphi_\xi) = \hat{I}(T(\varphi_\xi)) = \hat{I}(\xi).$$

It follows that $\tilde{I} = \hat{I}$.

Step 5: $\hat{I}$ is continuous. Suppose that I is continuous. Fix any $\xi, \xi' \in B_0(\mathcal{D}_\mathcal{M})$ and $c \in \mathbb{R}$. First we show that the set $\{\alpha \in [0,1] : \hat{I}(\alpha\xi + (1-\alpha)\xi') \leq c\}$ is closed. If it is empty, it is closed. If it is nonempty, take any sequence $(\alpha_n)_n \subseteq L$ such that $\alpha_n \to \alpha_0$. By Lemma B.9, we can pick $\varphi, \varphi' \in B_0(\mathcal{A})$ such that $\xi = T(\varphi)$ and $\xi' = T(\varphi')$. Moreover, we can pick we a finite partition $(E_i)_{i=1}^k$ and reals $(r_i)_{i=1}^k$, $(r_i')_{i=1}^k$ such that:

$$\varphi = \sum_{i=1}^k \chi_{E_i} r_i, \quad \varphi' = \sum_{i=1}^k \chi_{E_i} r_i'.$$

Fix any $m \in \mathcal{M}$. Then, there is a unique E_{j_m} such that $m(E_{j_m}) = 1$ and $m(E_i) = 0$ if $i \neq j_m$. Therefore, it follows that for all $n \in \mathbb{N}$,

$$I^m(\alpha_n\varphi + (1-\alpha_n)\varphi') = \alpha_n r_{j_m} + (1-\alpha_n)r'_{j_m} = \alpha_n I^m(\varphi) + (1-\alpha_n)I^m(\varphi'),$$
$$I^m(\alpha_0\varphi + (1-\alpha_0)\varphi') = \alpha_0 r_{j_m} + (1-\alpha_0)r'_{j_m} = \alpha_0 I^m(\varphi) + (1-\alpha_0)I^m(\varphi').$$

Since $m \in \mathcal{M}$ was arbitrarily chosen, it follows that:

$$\forall n \in \mathbb{N}, \quad \alpha_n\xi + (1-\alpha_n)\xi = \alpha_n T(\varphi) + (1-\alpha_n)T(\varphi') = T(\alpha_n\varphi + (1-\alpha_n)\varphi')$$
$$\alpha_0\xi + (1-\alpha_0)\xi = \alpha_0 T(\varphi) + (1-\alpha_0)T(\varphi') = T(\alpha_0\varphi + (1-\alpha_0)\varphi')$$

Therefore, by definition of $\hat{I}$ and continuity of I:

$$\begin{aligned} c &\geq \liminf_n \hat{I}(\alpha_n\xi + (1-\alpha_n)\xi') \\ &= \liminf_n I(\alpha_n\varphi + (1-\alpha_n)\varphi') \\ &= I(\alpha_0\varphi + (1-\alpha_0)\varphi') \\ &= \hat{I}(\alpha_0\xi + (1-\alpha_0)\xi') \end{aligned}$$

and, therefore, $\alpha_0 \in \{\alpha \in [0,1] : \hat{I}(\alpha\xi + (1-\alpha)\xi') \leq c\}$, showing that this set is closed. By a symmetric argument, we can show that $\{\alpha \in [0,1] : \hat{I}(\alpha\xi + (1-\alpha)\xi') \geq c\}$ is also closed. Siince this holds for all $\xi, \xi' \in B_0(\mathcal{D}_\mathcal{M})$ and $c \in \mathbb{R}$, and $\hat{I}$ is monotone by Step 3, Proposition 43 in Cerreia-Vioglio et al. (2011) implies that $\hat{I}$ is continuous.

Step 6: $\hat{I}$ is quasiconcave. Fix any $\alpha \in \mathbb{R}$. We show that the set $U_c = \{\xi \in B_0(\mathcal{D}_\mathcal{M}) : \xi \geq c\}$ is convex. If it is empty, this holds vacuously true. Suppose it is nonempty. Take $\xi_1, \xi_2 \in U_c$ and $\alpha \in [0,1]$. By Lemma B.9, we can pick $\varphi_1, \varphi_2 \in B_0(\mathcal{A})$ such that $\xi_1 = T(\varphi_1$ and $\xi_2 = T = \varphi_2$. Notice that $I(\varphi_1) = \hat{I}(\xi_1) \geq c$ and $I(\varphi_1) = \hat{I}(\xi_1) \geq c$. Since I is quasiconcave, it follows that $I(\alpha\varphi_1 + (1-\alpha)\varphi_2) \geq c$. Now, pick a partition $\{E_i\}_{i=1}^k \subseteq \mathcal{F}$ and profiles of scalars $(r_i^1)_{i=1}^k, (r_i^2)_{i=1}^k \subseteq \mathbb{R}$ such that $\varphi_1 = \sum_{i=1}^{k} \chi_{E_i} r_i^1$ and $\varphi_2 = \sum_{i=2}^{k} \chi_{E_i} r_i^2$. Fix $m \in \mathcal{M}$. Since the partition is in $\mathcal{A}$, there is a unique j_m such that $m(E_{j_m}) = 1$ and $m(E_i) = 0$ if $i \neq j_m$. Therefore,

$$I^m(\alpha\varphi_1+(1-\alpha)\varphi_2) = \alpha r^1_{j_m}+(1-\alpha)r^2_{j_m} = \alpha I^m(\varphi_1)+(1-\alpha)I^m(\varphi_2) = \alpha\xi_1(m)+(1-\alpha)\xi_2(m)$$

Therefore, we can conclude that $T(\alpha\varphi_1 + (1-\alpha)\varphi_2) = \alpha\xi_1 + (1-\alpha)\xi_2$. Then:

$$\hat{I}(\alpha\xi_1 + (1-\alpha)\xi_2) = I(\alpha\xi_1 + (1-\alpha)\xi_2) \geq c$$

and, therefore, $\alpha\xi_1+(1-\alpha)\xi_2 \in U_c$, showing convexity. Since c was arbitrarily chosen, we conclude that $\hat{I}$ is quasiconcave.

Step 7: $\hat{I}$ is monotone continuous Take $\xi, \xi' \in B_0(\mathcal{D}_\mathcal{M})$ and $k \in \mathbb{R}$, a monotone sequence $(D_n)_n \in D_\mathcal{M}$ such that $D_n \downarrow \emptyset$, and assume that $\hat{I}(\xi) > \hat{I}(\xi')$. Then, we can find $\varphi, \varphi' \in B_0(\mathcal{A})$ such that $\xi = T(\varphi)$ and $T(\varphi') = \xi'$. It follows that $I(\varphi) = \hat{I}(\xi) > \hat{I}(\xi') = I(\varphi')$. Let $E_n := \mathfrak{q}^{-1}(D_n) \in \mathcal{A}$ and notice that $E_n \downarrow \emptyset$. Therefore, there exists n_0 such that $I(kE_{n_0}\varphi) > I(\varphi')$. Since $E_{n_0} \in \mathcal{A}$, for all $m \in \mathcal{M}$, $m(E_{n_0}) \in \{0, 1\}$ and

$$m(E_{n_0}) = 1 \implies kE_{n_0}\varphi = k \text{ a.e. } [m] \implies I^m(kE_{n_0}\varphi) = I^m(k) = k$$

$$m(E_{n_0}) = 0 \implies kE_{n_0}\varphi = \varphi \text{ a.e. } [m] \implies I^m(kE_{n_0}\varphi) = I^m(\varphi) = \xi(m)$$

Moreover, notice that $m(E_{n_0}) = 1$ if and only if $m \in D_{n_0}$ and $m(E_{n_0}) = 0$ if and only if $m \notin D_{n_0}$. Therefore, $kD_{n_0}\xi = T(kE_{n_0}\varphi)$ and we can conclude that $\hat{I}(kD_{n_0}\xi) = I(kE_{n_0}\varphi) > I(\varphi') = \hat{I}(\xi')$ as we wanted to show.

(ii) implies (i).

Suppose there exists a normalized, monotone, and continuous functional $\hat{I} : B_0(\mathcal{M}, \mathcal{D}_\mathcal{M}) \to \mathbb{R}$ such that for all $\varphi \in B_0(\mathcal{A})$, $I(\varphi) = \hat{I}(T(\varphi))$.

Step 1: I is normalized.

Take $k \in \mathbb{R}$. Since $\hat{I}$ is normalized, we have that $\hat{I}(k) = k$. Moreover, $T(k)(m) = I^m(k) = k$ for all $m \in \mathcal{M}$. Therefore, $I(k) = \hat{I}(T(k)) = \hat{I}(k) = k$, showing that I is normalized.

Step 2: I is monotone.

Take $\varphi, \varphi' \in B_0(\mathcal{A})$ such that $\varphi \geq \varphi'$. For all $m \in \mathcal{M}$, I^m is monotone and, therefore, $T(\varphi)(m) = I^m(\varphi) \geq I^m(\varphi') = T(\varphi')(m)$. But, then, since $\hat{I}$ is monotone

$$I(\varphi) = \hat{I}(T(\varphi)) \geq \hat{I}(T(\varphi')) = I(\varphi'),$$

showing that I is monotone.

Step 3: If $\varphi, \varphi' \in B_0(\mathcal{A})$ and $I^m(\varphi) \geq I^m(\varphi')$ for all $m \in \mathcal{M}$, then $I(\varphi) \geq I(\varphi')$.

Take any two $\varphi, \varphi' \in B_0(\mathcal{A})$ and assume that $I^m(\varphi) \geq I^m(\varphi')$ for all $m \in \mathcal{M}$. Then, $T(\varphi) \geq T(\varphi')$ and, therefore, since $\hat{I}$ is monotone:

$$I(\varphi) = \hat{I}(T(\varphi)) \geq \hat{I}(T(\varphi')) = I(\varphi').$$

Step 4: I is continuous. Take a sequence $(\varphi_n)_n \subseteq B_0(\mathcal{A})$ such that $\varphi_n \to \varphi \in B_0(\mathcal{A})$ uniformly. Since for each $m \in \mathcal{M}$, I^m is Lipschitz continuous, it follows that for all m, $|I^m(\varphi_n) - I^m(\varphi)| \leq ||\varphi - \varphi_n||_\infty$ so that:

$$||T(\varphi_n) - T(\varphi)||_\infty \leq ||\varphi - \varphi_n||_\infty \to 0.$$

Thus, $T(\varphi_n)$ converges uniformly to $T(\varphi)$ and by Lemma B.9, $T(\varphi_n), T(\varphi) \in B_0(\mathcal{D}_\mathcal{M})$. Therefore, by continuity of $\hat{I}$, we have that:

$$I(\varphi_n) = \hat{I}(T(\varphi_n)) \to \hat{I}(T(\varphi)) = I(\varphi)$$

showing that I is continuous.

Step 5: I is quasiconcave. Suppose $\hat{I}$ is quasiconcave. Take $\varphi_1, \varphi_2 \in B_0(\mathcal{A})$ and $\alpha \in [0,1]$. Since I^m is concave, it follows that

$$I^m(\alpha\varphi_1 + (1-\alpha)\varphi_2) \geq \alpha I^m(\varphi_1) + (1-\alpha)I^m(\varphi_2)$$

for all $m \in \mathcal{M}$. Therefore, since $\hat{I}$ is monotone and quasiconcave,

$$\begin{aligned} I(\alpha\varphi_2 + (1-\alpha)\varphi_2) &= \hat{I}(T(\alpha\varphi_1 + (1-\alpha)\varphi_2)) \\ &\geq \hat{I}(\alpha T(\varphi_1) + (1-\alpha)T(\varphi_2)) \\ &\geq \min\{\hat{I}(T(\varphi_1)), \hat{I}(T(\varphi_2))\} = \min\{I(\varphi_1), I(\varphi_2)\} \end{aligned}$$

showing that I is quasiconcave.

Step 6: I is monotone continous. Take $\varphi, \varphi' \in B_0(\mathcal{A})$ and $k \in \mathbb{R}$, a monotone sequence $(E_n)_n \in \mathcal{A}$ such that $E_n \downarrow \emptyset$, and assume that $I(\varphi) > I(\varphi')$. Then, $\hat{I}(T(\varphi)) = I(\varphi) > I(\varphi') = \hat{I}(T(\varphi'))$. Notice that for each $n \in \mathbb{N}$, $E_n \in \mathcal{A}$ and, therefore, $m(E_n) \in \{0,1\}$ for all $m \in \mathcal{M}$. Then, let $D_n = \{m \in \mathcal{M} : m(E_n) > \frac{1}{2}\}$ and notice that $m \in D_n$ if and only if $m(E_n) = 1$ and $m \notin D_n$ if and only if $m(E_n) = 0$. Clearly, D_n is a decreasing sequence of sets. We show that $\cap_n D_n = \emptyset$. Take any $m \in \mathcal{M}$. Since m is countably additive, by continuity of finite measures, it must be the case that $m(E_n) \to 0$. However, since $m(E_n) \in \{0,1\}$ for all $n \in \mathbb{N}$, this implies that there is a N such that $m(E_n) = 0$ for all $n > N$. This implies that $m \notin E_n$ for $n > N$ and, therefore, $m \notin \cap_n D_n$. It follows that $D_n \downarrow \emptyset$. Since $\hat{I}$ is monotone continuous, there exists a n_0 such that $\hat{I}(\chi_{D_{n_0}} k + \chi_{D^c_{n_0}} T(\varphi)) > \hat{I}(T(\varphi'))$. Finally note that for all $m \in \mathcal{M}$,

$$\begin{aligned} m \in D_{n_0} &\implies m(D_{n_0}) = 1 \implies I^m(\chi_{D_{n_0}} k + \chi_{D^c_{n_0}} \varphi) = I^m(k) = k \\ m \in D^c_{n_0} &\implies m(D_{n_0}) = 0 \implies I^m(\chi_{D_{n_0}} k + \chi_{D^c_{n_0}} \varphi) = I^m(\varphi) = T(\varphi)(m). \end{aligned}$$

Hence, $T(\chi_{D_{n_0}}k+\chi_{D^c_{n_0}}\varphi)=\chi_{D_{n_0}}k+\chi_{D^c_{n_0}}T(\varphi)$ and, therefore,

$$\begin{aligned} I(\chi_{D_{n_0}}k+\chi_{D^c_{n_0}}\varphi) &= \hat{I}(T((\chi_{D_{n_0}}k+\chi_{D^c_{n_0}}\varphi)) \\ &= \hat{I}(\chi_{D_{n_0}}k+\chi_{D^c_{n_0}}T(\varphi)) \\ &> \hat{I}(T(\varphi')) = I(\varphi') \end{aligned}$$

as we wanted to show. ■

Appendix E. Proof of Theorem 4

Proof of Theorem 4: We know that $\succsim$ is represented by u when restricted to constant acts. Define the functional $I : B_0(\mathcal{G}) \to \mathbb{R}$ such that for each $\varphi \in B_0(\mathcal{G})$, $I(\varphi) \coloneqq u(x_{f_\varphi})$, where $f_\varphi \in \mathcal{F}$ is chosen so that $\varphi = u(f_\varphi)$. By Lemma B.2, such act f_φ exists for all $\varphi \in B_0(\mathcal{G})$, while the certainty equivalent $x_{f_\varphi} \sim f_\varphi$ exists by Lemma B.3. Moreover, for any $\varphi \in B_0(\mathcal{G})$, if there are two $f_\varphi, f'_\varphi \in \mathcal{F}$ such that $u(f_\varphi) = \varphi = u(f'_\varphi)$, we then have that since u represents $\succsim$ over X,

$$\begin{aligned} u(f_\varphi)(\omega) = u(f'_\varphi)(\omega) &\implies u(f_\varphi(\omega)) = u(f'_\varphi(\omega)) \\ &\implies f_\varphi(\omega) \sim f'_\varphi(\omega) \end{aligned}$$

for all $\omega \in \Omega$. By Axiom 1.(ii) of monotonicity, it follows that $f_\varphi \sim f'_\varphi$ and, by transitivity, that $x_{f_\varphi} \sim x_{f'_\varphi}$. Therefore, we can conclude that $u(x_{f_\varphi}) = u(x_{f'_\varphi})$, showing that I is a well-defined functional on $B_0(\mathcal{G})$. It is easily seen that such functional is also normalized, monotone, and continuous.[14] Moreover, it is monotone continuous and its restriction to $B_0(\mathcal{A})$ is quasiconcave.

[14] See for example the proof of Theorem 1 (Omnibus) in the working paper version of Cerreia-Vioglio et al. (2022).

Define the function $V \coloneqq I \circ u : \mathcal{F} \to \mathbb{R}$. For all $f, f' \in \mathcal{F}$,

$$\begin{aligned} f \succsim f' &\iff x_{f'} \succsim x_{f'} \\ &\iff V(f) = I(u(f)) = u(x_f) \geq u(x_{f'}) = I(u(f)) = V(f') \ . \end{aligned}$$

This shows that V represents $\succsim$ on $\mathcal{F}$. Moreover, by Proposition 1, for each $m \in \mathcal{M}$, $\succsim^m$ is represented by $I^m \circ u$, where $I^m : B(\mathcal{G}) \to \mathbb{R}$ is as defined in (12). Moreover, let $I_{\mathcal{A}}$ be the generalized conditional expectation as in Corollary 3. Take now $\varphi, \psi \in B_0(\mathcal{G})$ such that $I^m(\varphi) \geq I^m(\psi)$ for all $m \in \mathcal{M}$. By Lemma B.2, we can find $f_\varphi, f_\psi \in \mathcal{F}$ such that $\varphi = u(f_\varphi)$ and $\psi = u(f_\psi)$. Then, $I^m(u(f_\varphi)) \geq I^m(u(f_\psi))$ for all $m \in \mathcal{M}$ so that $f_\varphi \succsim^m f_\psi$ for all $m \in \mathcal{M}$. Consistency implies that $f_\varphi \succsim f_\psi$. Therefore:

$$I(\varphi) = I(u(f_\varphi)) = V(f_\varphi) \geq V(f_\psi) = I(u(f_\psi)) \geq I(\psi) \ .$$

By this fact and since I is monotone, normalized, continuous, and quasiconcave, by Lemma B.10, there exists a unique monotone, normalized, continuous, and quasiconcave functional $\hat{I} : B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ such that $I(\varphi) = \hat{I}(T(\varphi))$ for all $\varphi \in B_0(\mathcal{A})$. Moreover, since I is monotone continuous, so is $\hat{I}$. By Theorem 21 in Cerreia-Vioglio et al. (2013), $\hat{I}$ admits a unique monotone, normalized, lower semicontinuous, and quasiconcave extension to $B(\mathcal{D}_{\mathcal{M}})$, which, abusing notation, we also denote by $\hat{I}$. Moreover, since $\hat{I}$ is monotone continuous when restricted to $B_0(\mathcal{D}_{\mathcal{M}})$, it is inner/outer continuous on $B(\mathcal{D}_{\mathcal{M}})$. Take now any $\varphi \in B_0(\mathcal{G})$. Since $I_{\mathcal{A}}(\varphi) \in B(\mathcal{A})$ and $B_0(\mathcal{A})$ is dense in $B(\mathcal{A})$, we can pick sequences $(\psi_n^l)_{n \in \mathbb{N}}, (\psi_n^u)_{n \in \mathbb{N}} \in B_0(\mathcal{A})$ such that $\psi_n^l \nearrow I_{\mathcal{A}}(\varphi)$ and $\psi_n^u \searrow I_{\mathcal{A}}(\varphi)$ uniformly. Fix any $m \in \mathcal{M}$. Since I^m is monotone, we have that for all $n \in \mathbb{N}$:

$$I^m(\psi_n^l) \leq I^m(I_{\mathcal{A}}(\varphi)) \leq I^m(\psi_n^u) \ .$$

By Proposition 1, we also have that $I^m(I_{\mathcal{A}}(\varphi)) = I^m(\varphi)$ and, therefore, we have that for all $n \in \mathbb{N}$,

$$I^m(\psi_n^l) \leq I^m(\varphi) \leq I^m(\psi_n^u)$$

for all $m \in \mathcal{M}$. By what shown above, we have that for all $n \in \mathbb{N}$:

$$\hat{I}(T(\psi_n^l)) = I(\psi_n^l) \leq I(\varphi) \leq I(\psi_n^u) \leq I(T(\psi_u^l)).$$

Since I^m is monotone and Lipschitz, we have that $T(\psi_n^l) \uparrow T(I_{\mathcal{A}}(\varphi)) = T(\varphi)$ and $T(\psi_n^u) \downarrow T(I_{\mathcal{A}}(\varphi)) = T(\varphi)$. Then, since $\hat{I}$ is inner/outer continuous, passing to the limit in the above sequence of inequality, we obtain:

$$\hat{I}(T(\varphi)) = \lim_n \hat{I}(T(\psi_n^l)) \leq I(\varphi) \leq \lim \hat{I}(T(\psi_u^l)) = \hat{I}(T(\varphi)).$$

This shows that $I(\varphi) = \hat{I}(T(\varphi))$ for all $\varphi \in B_0(\mathcal{G})$. It follows that for all $f, g \in \mathcal{F}$,

$$f \succsim g \iff I(u(f)) \geq I(u(g)) \iff \hat{I}(T(u(f))) \geq \hat{I}(T(u(g))).$$

■

Proof of Proposition 5: Let $\succsim_1$ and $\succsim_2$ Suppose that $\succsim_1$ and $\succsim_2$ are two misspecification averse preferences represented respectively by $(\hat{I}_1, u_1, c_1)$ and $(\hat{I}_2, u_2, c_2)$ as in Theorem 4. Suppose that $u_1 = u_2 = u$ and that $\hat{I}_1 \leq \hat{I}_2$. Take any $f \in \mathcal{F}(\mathcal{A})$ and $x \in X$ and assume that $f \succsim_1 x$. Since f is measurable with respect to $\mathcal{A}$, for each $m \in \mathcal{M}$, f must be constant on E^m and, therefore coherence and normalziation imply:

$$I_1(u(f), m) = I_1\left(u(f)\chi_{E^m}, m\right) = u(f|_{E^m}) = I_2\left(u(f)\chi_{E^m}, m\right) = I_2(u(f), m).$$

Then, we have that:

$$u(x) \leq \hat{I}_1\left(I_1(u(f), \cdot)\right) \leq \hat{I}_2\left(I_1(u(f), \cdot)\right) = \hat{I}_2\left(I_2(u(f), \cdot)\right)$$

so that $f \succsim_2 x$.

As for the other direction, equation (9) and nontriviality automatically imply that u_2 is a positive affine transformation of u_1. Assume that $u_1 = u_2 = u$ and take $\xi \in B_0(\mathcal{D}_{\mathcal{M}})$. Then, by Lemmas B.9 and B.2, there exists $f \in \mathcal{F}(\mathcal{A})$ such that $\xi = I_1(u(f), \cdot)$. By the same argument given above, it is also the case that $\xi = I_2(u(f), \cdot)$. Take $x \in X$ such that $f \sim_1 x$. Then, condition (9) implies that $f \succsim_2 x$. Therefore:

$$\hat{I}_1(\xi) = \hat{I}_1(I_1(u(f), \cdot)) = u(x) \leq \hat{I}_2(I_2(u(f), \cdot)) = \hat{I}_2(\xi).$$

Thus, $\hat{I}_1 \leq \hat{I}_2$ on $B_0(\mathcal{D}_{\mathcal{M}})$. Since this set is dense in $B(\mathcal{D}_{\mathcal{M}})$, we can find a monotonically decreasing sequence $(\xi_n)_n \subseteq B_0(\mathcal{D}_{\mathcal{M}})$ such that $\xi_n \searrow \xi$. Then, $\hat{I}_1(\xi_n) \leq \hat{I}_2(\xi_n)$ for all $n \in \mathbb{N}$. Since $\hat{I}$ is inner/outer continuous, passing to the limit we can conclude that $\hat{I}_1(\xi) \leq \hat{I}_2(\xi)$. ■

E.1. **Proof of Theorems 6 and 7.** In this section we prove the general representation in Theorem 6. We start with the following lemma.

LEMMA B.12: *Suppose $\succsim$ is a misspecification averse preference whose restriction to $\mathcal{F}(\mathcal{A})$ satisfies Savage's P2-P6. There exist a non-constant, affine, and surjective $\tilde{u} : X \to \mathbb{R}$, a strictly increasing $\phi : \mathbb{R} \to \mathbb{R}$, and a non-atomic $\nu \in \Delta^\sigma(\Omega, \mathcal{A})$ such that for all $f, g \in \mathcal{F}(\mathcal{A})$,*

$$f \succsim g \iff \phi^{-1}\left(\int_\Omega \phi(\tilde{u}(f))d\nu\right) \geq \phi^{-1}\left(\int_\Omega \phi(\tilde{u}(g))d\nu\right) .$$

Moreover, ν is unique, $\tilde{u}$ is unique up to positive affine transformations, and ϕ is unique up to positive affine transformations given $\tilde{u}$.

Proof of Lemma B.12: Since when restricted to acts measurable with respect to $\mathcal{A}$, $\succsim$ satisfies the Axioms of Savage (1954) and monotone continuity, there exist a non-constant function $v : X \to \mathbb{R}$ and a non-atomic probability measure $\nu \in \Delta^\sigma(\Omega, \mathcal{A})$ such that for all $f, f' \in \mathcal{F}(\mathcal{A})$:

$$f \succsim f' \iff \int_\Omega v(f) d\nu \geq \int_\Omega v(f') d\nu.$$

Clearly, v represents $\succsim$ on X. By Herstein and Milnor (1953), there exists an affine $\tilde{u} : X \to \mathbb{R}$ representing $\succsim$ on X. Since $\succsim$ is unbounded, the argument in the proof of Proposition 1 shows that u must be surjective. Then, there exists a strictly increasing trasformation $\phi : \mathbb{R} \to \mathbb{R}$ such that $v = \phi \circ u$. Now, Take any $k, k' \in \operatorname{Im} \phi$ and $\lambda \in (0, 1)$. Then, we can find $x_k, x_{k'} \in X$ such that $\phi(k) = \phi(u(x_k))$ and $\phi(k') = \phi(u(x_{k'}))$. Since ν is non-atomic, we can pick $E_\lambda \in \mathcal{A}$ such that $\nu(E_\lambda) = \lambda$. Then, $f_\lambda := x_k E_\lambda x_{k'} \in \mathcal{F}(\mathcal{A})$ and, by Lemma B.3, we can find $x_{f_\lambda} \in X$ such that $x_{f_\lambda} \sim f_\lambda$. Clearly, both f_λ and x_{f_λ} are measurable with respect to $\mathcal{A}$. Therefore:

$$\begin{aligned}
\lambda\phi(k) + (1-\lambda)\phi(k'), &= \nu(E_\lambda)\phi(u(x_k)) + \nu(\Omega \setminus E_\lambda)\phi(u(x_{k'})) \\
&= \int_\Omega \phi(u(f_\lambda)) d\nu \\
&= \int_\Omega \phi(u(x_{f_\lambda})) d\nu \\
&= \phi(u(x_{f_\lambda})) \in \operatorname{Im} \phi \ .
\end{aligned}$$

Thus, $\phi : \mathbb{R} \to \mathbb{R}$ is strictly increasing and has a convex image. It follows that ϕ is continuous.

The uniqueness of the representation follows by standard arguments. ■

Lemma B.13: *Suppose $(\Omega, \mathcal{G}, \mathcal{M})$ is a structured space and there exist a utility function $u : X \to \mathbb{R}$, a convex statistical distance $c : \Delta \times \mathcal{M} \to [0, \infty]$, a strictly increasing and continuous function $\phi : \operatorname{Im} u \to \mathbb{R}$ and a prior $\mu \in \Delta^\sigma(\mathcal{M}, \mathcal{D}_\mathcal{M})$ such* *that* $\succsim$ *is represented on* $\mathcal{F}$ *by*

$$\int_{\mathcal{M}} \phi\left(I^m(u(f))\right) d\mu(m)$$

where I^m *is defined as in* (5). *Then, there exists a probability measure* $\nu \in \Delta^\sigma(\Omega, \mathcal{A})$ *such that the restriction of* $\succsim$ *to* $\mathcal{F}(\mathcal{A})$ *is represented by*

$$\int_{\Omega} \phi(u(f)) d\nu.$$

Moreover, ν *is nonatomic if* μ *is nonatomic.*

Proof of Lemma B.13: Suppose the premise holds and define the following measure: for all $A \in \mathcal{A}$,

$$\nu(A) = \int_{\mathcal{M}} m(A) d\mu(m)$$

and notice that $\nu \in \Delta^\sigma(\Omega, \mathcal{A})$ and $\nu(\Omega_0) = 1$. Moreover, for all $D \in \mathcal{D}_{\mathcal{M}}$, since

$$m \in D \implies m\left(\{\omega \in \Omega : \mathfrak{q}(\omega) \in D\}\right) \geq m\left(\{\omega \in \Omega : \mathfrak{q}(\omega) = m\}\right) = 1,$$

$$m \notin D \implies m\left(\{\omega \in \Omega : \mathfrak{q}(\omega) \in D\}\right) \leq 1 - m\left(\{\omega \in \Omega : \mathfrak{q}(\omega) = m\}\right) = 0$$

then,

$$\begin{aligned}
\nu \circ \mathfrak{q}^{-1}(D) &= \nu\left(\{\omega \in \Omega : \mathfrak{q}(\omega) \in D\}\right) \\
&= \int_{\mathcal{M}} m(\{\omega \in \Omega : \mathfrak{q}(\omega) \in D\})\, d\mu(m) \\
&= \int_{D} m(\{\omega \in \Omega : \mathfrak{q}(\omega) \in D\})\, d\mu(m) + \int_{\mathcal{M}\setminus D} m(\{\omega \in \Omega : \mathfrak{q}(\omega) \in D\})\, d\mu(m) \\
&= \int_{D} 1 d\mu(m) = \mu(D).
\end{aligned}$$

Therefore, for any $\psi \in B_0(\Omega, \mathcal{A})$, we have that

$$\begin{aligned}\int_{\mathcal{M}} \phi\left(I^m(\psi)\right) d\mu(m) &= \int_{\mathcal{M}} \phi\left(I^m(\psi)\right) d(\nu \circ \mathfrak{q}^{-1})(m) \\ &= \int_{\Omega_0} \phi\left(I^{\mathfrak{q}(\omega)}(\psi)\right) d\nu(\omega) \\ &= \int_{\Omega} \phi\left(I_{\mathcal{A}}(\psi)(\omega)\right) d\nu(\omega) \\ &= \int_{\Omega} \phi\left(\psi\right) d\nu.\end{aligned}$$

where we apply the change of variable formula and $I_{\mathcal{A}}$ is the generalized common conditional expectation of $(I^m)_{m\in\mathcal{M}}$ given $\mathcal{A}$ of Corollary 3. It follows that for all $f, g \in \mathcal{A}$, $f \succsim f$ if and only if

$$\int_{\Omega} \phi(u(f))d\nu \geq \int_{\Omega} \phi(u(g))d\nu.$$

as we wanted to show.

Furthermore, assume that μ is notatomic. We show that also ν is non-atomic. To this end, take $E \in \mathcal{A}$ such that $\nu(A) > 0$. Then, there exists by Lemma B.9 a set $D_E \in \mathcal{D}_{\mathcal{M}}$ such that $I^m(\chi_E) = \chi_{D_E}(m)$ for all $m \in \mathcal{M}$. Then,

$$\mu(D_E) = \int_{\mathcal{M}} \chi_{D_E}(m)d\mu(m) = \int_{\mathcal{M}} I^m(\chi_E)d\mu(m) = \int_{\mathcal{M}} m(E)d\mu(m) = \nu(E) > 0$$

where we use the fact that $m(E) \in \{0, 1\}$ for all $m \in \mathcal{M}$. Since μ is nonatomic, there exists a subset $D_0 \subseteq D_E$ in $\mathcal{D}_{\mathcal{M}}$ such that $0 < \mu(D_0) < \mu(D_E)$. Again by Lemma B.9, we can find $E^{D_0} \in \mathcal{A}$ such that $\chi_{D_0}(m) = I^m(\chi_{E^{D_0}})$ for all $m \in \mathcal{M}$. Then, let $E_0 \coloneqq E \cap E^{D_0} \subseteq E$. We have that for all $m \in \mathcal{M}$,

$$\chi_{D_0} = \chi_{D_0}\ \chi_{D_E} = I^m(\chi_{E^{D_0}})\ I^m(\chi_E) = I^m(\chi_{E_0})$$

where again we use Lemma B.1. Therefore,

$$\nu(E_0) = \int_{\mathcal{M}} m(E_0) d\mu(m) = \int_{\mathcal{M}} I^m(\chi_{E_0}) d\mu(m) = \int_{\mathcal{M}} \chi_{D_0} d\mu(m) = \mu(D_0)$$

so that $0 < \nu(E_0) < \nu(E)$, proving that ν is nonatomic. ■

Proof of Theorem 6: *(i) implies (ii).*

We know that $\succsim$ is represented by u when restricted to constant acts. Define the functional $I : B_0(\mathcal{G}) \to \mathbb{R}$ such that for each $\varphi \in B_0(\mathcal{G})$, $I(\varphi) \coloneqq u(x_{f_\varphi})$, where $f_\varphi \in \mathcal{F}$ is chosen so that $\varphi = u(f_\varphi)$. By Lemma B.2, such act f_φ exists for all $\varphi \in B_0(\mathcal{G})$, while the certainty equivalent $x_{f_\varphi} \sim f_\varphi$ exists by Lemma B.3. Moreover, for any $\varphi \in B_0(\mathcal{G})$, if there are two $f_\varphi, f'_\varphi \in \mathcal{F}$ such that $u(f_\varphi) = \varphi = u(f'_\varphi)$, we then have that since u represents $\succsim$ over X,

$$\begin{aligned} u(f_\varphi)(\omega) = u(f'_\varphi)(\omega) &\implies u(f_\varphi(\omega)) = u(f'_\varphi(\omega)) \\ &\implies f_\varphi(\omega) \sim f'_\varphi(\omega) \end{aligned}$$

for all $\omega \in \Omega$. By Axiom 1.(ii) of monotonicity, it follows that $f_\varphi \sim f'_\varphi$ and, by transitivity, that $x_{f_\varphi} \sim x_{f'_\varphi}$. Therefore, we can conclude that

$$I(\varphi) = u(x_{f_\varphi}) = u(x_{f'_\varphi}) = I(f'_\varphi)$$

showing that I is a well-defined functional on $B_0(\mathcal{G})$. It is easily seen that such functional is also normalized, monotone, and continuous.[15]

Define the function $V \coloneqq I \circ u : \mathcal{F} \to \mathbb{R}$. For all $f, f' \in \mathcal{F}$,

$$\begin{aligned} f \succsim f' &\iff x_{f'} \succsim x_{f'} \\ &\iff V(f) = I(u(f)) = u(x_f) \geq u(x_{f'}) = I(u(f)) = V(f') \ . \end{aligned}$$

[15]See for example the proof of Theorem 1 (Omnibus) in the working paper version of Cerreia-Vioglio et al. (2022).

This shows that V represents $\succsim$ on $\mathcal{F}$. Moreover, by Proposition 1, for each $m \in \mathcal{M}$, $\succsim^m$ is represented by $I^m \circ u$, where $I^m : B(\mathcal{G}) \to \mathbb{R}$ is as defined in (12). Moreover, let $I_{\mathcal{A}}$ be the generalized conditional expectation from Corollary 3. Take now $\varphi, \psi \in B_0(\mathcal{G})$ such that $I^m(\varphi) \geq I^m(\psi)$ for all $m \in \mathcal{M}$. By Lemma B.2, we can find $f_\varphi, f_\psi \in \mathcal{F}$ such that $\varphi = u(f_\varphi)$ and $\psi = u(f_\psi)$. Then, $I^m(u(f_\varphi)) \geq I^m(u(f_\psi))$ for all $m \in \mathcal{M}$ so that $f_\varphi \succsim^m f_\psi$ for all $m \in \mathcal{M}$. Consistency implies that $f_\varphi \succsim f_\psi$. Therefore:

$$I(\varphi) = I(u(f_\varphi)) = V(f_\varphi) \geq V(f_\psi) = I(u(f_\psi)) \geq I(\psi) \ .$$

Moreover, by Lemma B.12, there exist an unbounded and affine $\tilde{u} : X \to \mathbb{R}$, a strictly increasing $\phi : \mathbb{R} \to \mathbb{R}$, and a non-atomic $\nu \in \Delta^\sigma(\Omega, \mathcal{A})$ such that the restriction of $\succsim$ to $\mathcal{F}(\mathcal{A})$ is represented by the functional:

$$f \mapsto \phi^{-1}\left(\int_\Omega \phi(\tilde{u}(f))d\nu\right) \ .$$

Moreover, since $\Omega \setminus \Omega_0$ is null, $\nu(\Omega \setminus \Omega_0) = 0$. Without loss of generality, we can assume that $\tilde{u} = u$ and normalize $\phi(0) = 0$ and $\phi(1) = 1$. Now, define the map $J : B(\mathcal{A}) \to \mathbb{R}$ such that

$$J(\varphi) = \phi^{-1}\left(\int_\Omega \phi(\varphi)d\nu\right)$$

for all $\varphi \in B(\mathcal{A})$. Since ϕ is continuous and strictly increasing, J is well-defined, normalized, and continuous. Moreover, for all $f, g \in \mathcal{F}(\mathcal{A})$,

$$f \succsim g \iff J(u(f)) \geq J(u(g)) \ .$$

Moreover, take any $\varphi \in B_0(\mathcal{A})$. By Lemma B.2, we can choose $f_\varphi \in \mathcal{F}(\mathcal{A})$ such that $\varphi = u(f_\varphi) = f_\varphi$. Then, since both V and $J \circ u$ represent $\succsim$ on $\mathcal{F}(\mathcal{A})$,

$$I(\varphi) = I(u(f_\varphi)) = V(f_\varphi) = u(x_{f_\varphi}) = J\left(u(f_\varphi)\right) = J(\varphi) \ .$$

We conclude that $I(\varphi) = J(\varphi)$ for all $\varphi \in B_0(\mathcal{A})$. Take now any $\varphi \in B_0(\mathcal{G})$. Since $I_{\mathcal{A}}(\varphi) \in B(\mathcal{A})$ and $B_0(\mathcal{A})$ is dense in $B(\mathcal{A})$, we can pick sequences $(\psi_n^l)_{n\in\mathbb{N}}, (\psi_n^u)_{n\in\mathbb{N}} \in B_0(\mathcal{A})$ such that $\psi_n^l \nearrow I_{\mathcal{A}}(\varphi)$ and $\psi_n^u \searrow I_{\mathcal{A}}(\varphi)$ uniformly. Fix any $m \in \mathcal{M}$. Since I^m is monotone, we have that for all $n \in \mathbb{N}$:

$$I^m(\psi_n^l) \leq I^m(I_{\mathcal{A}}(\varphi)) \leq I^m(\psi_n^u) \ .$$

By Proposition 1, we also have that $I^m(I_{\mathcal{A}}(\varphi)) = I^m(\varphi)$ and, therefore, we have that for all $n \in \mathbb{N}$,

$$I^m(\psi_n^l) \leq I^m(\varphi) \leq I^m(\psi_n^u) \ .$$

Since m was chosen arbitrarily, this holds for all $m \in \mathcal{M}$. This and the fact that I and J coincide on $B_0(\mathcal{A})$ imply that for all $n \in \mathbb{N}$:

$$J(\psi_n^l) = I(\psi_n^l) \leq I(\varphi) \leq I(\psi_n^u) = J(\psi_n^u)$$

Passing to the limit and using the fact that J is continuous, we obtain that:

$$J(I_{\mathcal{A}}(\varphi)) \leq I(\varphi) \leq J(I_{\mathcal{A}}(\varphi)) \ .$$

That is:

$$\begin{aligned} I(\varphi) &= J(I_{\mathcal{A}}(\varphi)) \\ &= \phi^{-1}\left(\int_{\Omega} \phi\left(I_{\mathcal{A}}(\varphi)\right) \nu(d\tilde{\omega})\right) \\ &= \phi^{-1}\left(\int_{\Omega_0} \phi\left(\min_{p\in\Delta}\left\{\int_{\Omega} \varphi \ dp + c(p, \mathfrak{q}(\tilde{\omega}))\right\}\right) \nu(d\tilde{\omega})\right) \ . \end{aligned}$$

Finally, since $\mathfrak{q}_0 = \mathfrak{q}|_{\Omega_0}$ is a measurable transformation from $(\Omega_0, \mathcal{A}_0)$ to $(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$, define the image measure $\mu \coloneqq \nu \circ \mathfrak{q}_0^{-1} \in \Delta^{\sigma}(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$. Then, by Theorem 16.23 in

Billingsley (1995):

$$\begin{aligned} I(\varphi) &= \phi^{-1}\left(\int_{\Omega_0} \phi\left(\min_{p\in\Delta}\left\{\int_\Omega \varphi\ dp + c(p,\mathfrak{q}(\tilde\omega))\right\}\right) d\nu(\tilde\omega)\right) \\ &= \phi^{-1}\left(\int_{\mathcal{M}} \phi\left(\min_{p\in\Delta}\left\{\int_\Omega \varphi\ dp + c(p,\mathfrak{q}(\tilde\omega))\right\}\right) d(\nu\circ\mathfrak{q}^{-1})(m)\right) \\ &= \phi^{-1}\left(\int_{\mathcal{M}} \phi\left(\min_{p\in\Delta}\left\{\int_\Omega \varphi\ dp + c(p,m)\right\}\right) d\mu(m)\right) \ . \end{aligned}$$

But, then, $\succsim$ is represented on $\mathcal{F}$ by

$$V(f) = I(u(f)) = \phi^{-1}\left(\int_{\mathcal{M}} \phi\left(\min_{p\in\Delta}\left\{\int_\Omega u(f)\ dp + c(p,m)\right\}\right) d\mu(m)\right)$$

as we wanted to show. Next, we show that if $\chi_D = T(\chi_E)$ for $E \in \mathcal{A}$ and $D \in \mathcal{D}_{\mathcal{M}}$, then $\nu(E) = \mu(D)$. Indeed,

$$\begin{aligned} \phi^{-1}(\nu(E)) &= \phi^{-1}\left(\int_{\Omega_0} \phi(\chi_E) d\nu\right) \\ &= J(\chi_{E^D}) = J(I_{\mathcal{A}}(\chi_E)) = I(\chi_E) \\ &= \phi^{-1}\left(\int_{\mathcal{M}} \phi\left(I(\chi_E, m)\right) d\mu(m)\right) \\ &= \phi^{-1}\left(\int_{\mathcal{M}} \phi\left(\chi_D\right) d\mu(m)\right) \\ &= \phi^{-1}(\mu(D)). \end{aligned}$$

and since ϕ^{-1} is strictly increasing, this implies that $\nu(E) = \mu(D)$. We now show that μ is nonatomic. Take $D \in \mathcal{D}_{\mathcal{M}}$ such that $\mu(D) > 0$. By Lemma B.9, there exists $E^D \in \mathcal{A}$ such that $\chi_D = T(\chi_{E^D})$. Therefore, by what shown above, $\nu(E^D) = \mu(D) > 0$. Since ν is nonatomic, we can find $E_0 \in \mathcal{A}$ such that $E_0 \subseteq E^D$ and $\nu E_0 > 0$. By Lemma B.9, we can find $D_{E_0} \in \mathcal{D}_{\mathcal{M}}$ such that $\chi_{D_{E_0}} = T(\chi_{E_0})$. Suppose $m \in D_{E_0}$. Then, $I(\chi_{E_0}, m) = 1$ and since $E_0 \in \mathcal{A}$, it must be the case that $m(E_0) = 1$. Since $E_0 \subseteq E^D$, $m(E^D) = 1$ and, therefore, $\chi_D(m) = I(\chi_{E^D}, m) = 1$, so that $m \in D$. This shows that $D_{E_0} \subseteq D$. Moreover, by what shows above $\mu(D_{E_0}) = \nu(E_0) > 0$.

This shows that μ is non-atomic. It only remains to show that ϕ is concave. Take $r_1, r_2 \in \mathbb{R}$ and $\alpha = 1/2$. Since ν is nonatomic, we can find E such that $\nu(E) = 1/2$. Moreover, we can pick $x_1, x_2 \in X$ such that $r_1 = u(x_1)$ and $r_2 = u(x_2)$. Then:

$$\begin{aligned} J(u(x_1 E x_2)) &= \phi^{-1}\left(\int_\Omega \phi\left(u(x_1 E x_2)\right) d\nu\right) \\ &= \phi^{-1}\left(\frac{1}{2}\phi(r_1) + \frac{1}{2}\phi(r_2)\right) \\ &= \phi^{-1}\left(\int_\Omega \phi\left(u(x_2 E x_1)\right) d\nu\right) = J(u(x_2 E x_1)). \end{aligned}$$

Thus, $x_1 E x_2 \sim x_2 E x_1$. Since $\succsim$ satisfies uncertainty aversion, it follows that

$$\frac{1}{2}x_1 + \frac{1}{2}x_2 = \frac{1}{2}x_1 E x_2 + \frac{1}{2}x_2 E x_1 \succsim x_1 E x_2$$

and, therefore, since ϕ is increasing:

$$\begin{aligned} \phi\left(\frac{1}{2}r_1 + \frac{1}{2}r_2\right) &= \phi\left(J\left(u\left(\frac{1}{2}x_1 + \frac{1}{2}x_2\right)\right)\right) \\ &\geq \phi\left(J\left(u\left(x_1 E x_2\right)\right)\right) \\ &= \frac{1}{2}\phi(r_1) + \frac{1}{2}\phi(r_2). \end{aligned}$$

This shows that ϕ is midpoint concave. Since it is also strictly increasing on the interval $\mathbb{R}$, we conclude it is concave.

Uniqueness follows by standard arguments.

(ii) implies (i). It is clear that $\succsim$ satisfies Axioms 1-5 are satisfied. Moreover, by Lemma B.13, there exists a nonatomic probability measure $\nu \in \Delta^\sigma(\Omega, \mathcal{A}$ such that the restriction of $\succsim$ to $\mathcal{F}(\mathcal{A})$ is represented by the functional

$$\int_\Omega \phi(u(f))d\nu.$$

This implies that $\succsim$ satisfies Savage (1954)'s P2-P6 when restricted to $\mathcal{F}(\mathcal{A})$. ■

Proof of Theorem 7: *(i) implies (ii).* We know that there exists an affine $u : X \to \mathbb{R}$ and a normalized, monotone, continuous, and quasiconcave functional $I : B_0(\mathcal{G}) \to \mathbb{R}$ such that $\succsim$ is represented by $I \circ u$ on $\mathcal{F}$. By Proposition 1, we know that for each $m \in \mathcal{M}$, there exists I^m given as in (5) such that $I^m \circ u$ represents $\succsim^m$ on $\mathcal{F}$. By consistency, we also know that for all $\varphi, \psi \in B_0(\mathcal{G})$, $I^m(\varphi) \geq I^m(\psi)$ for all $m \in \mathcal{M}$ implies that $I(\varphi) \geq I(\psi)$. Therefore, by Proposition B.11, there exists a unique normalized, monotone, and continuous $\hat{I} : B_0(\mathcal{D}_{\mathcal{M}})$ such that $\hat{I}(I(\varphi, \cdot)) = I(\varphi)$ for all $\varphi \in B_0(\mathcal{G})$. Moreover, $\hat{I}$ is quasiconcave and monotone continuous. Take $\xi \in B_0(\mathcal{D}_{\mathcal{M}})$. By Lemma B.9, we can find a $\varphi \in B_0(\mathcal{A})$ such that $\xi = T(\varphi)$ and $f \in B_0(\mathcal{G})$ such that $\varphi = u(f)$. Notice that since there exists a K such that $\xi(m) \geq K$ for all $m \in \mathcal{M}$, $r_0 \coloneqq \inf_{m \in \mathcal{M}} \xi(m) \geq K$ and, therefore, $r_0 \in \mathbb{R}$. Pick $r > r_0$. Then, we can find $x_0, x \in X$ such that $r_0 = u(x_0)$ and $r = u(x)$. Take a sequence $(\alpha_n) \in (0,1)$ such that $\alpha_n \downarrow 0$ and let $x_n = \alpha_n x + (1-\alpha_n)x_0$. Fix any $n \in \mathbb{N}$. By affinity of u,

$$u(x_n) = \alpha_n u(x) + (1-\alpha_n)u(x_0) = \alpha_n r + (1-\alpha_n)r_0 > r_0 = \inf_{m \in \mathcal{M}} \xi(m) = \inf_{m \in \mathcal{M}} I(u(f), m).$$

Therefore, there exists $m_n \in \mathcal{M}$ such that $u(x_n) > I(u(f), m_n)$. This implies that $x_n \succ_m f$ and, therefore, Caution implies that $x_n \succsim f$. That is,

$$\alpha_n r + (1-\alpha_n)r_0 = u(x_n) \geq I(u(f)) = I(\varphi) = \hat{I}(\xi).$$

This holds for all $n \in \mathbb{N}$ and passing to the limit, we obtain $r_0 \geq \hat{I}(\xi)$. On the other hand, we have that for all $m \in \mathcal{M}$, $r_0 = \inf_{m' \in \mathcal{M}} \xi(m') \leq \xi(m)$ and, therefore, since $\hat{I}$ is normalized and monotone, $r_0 = \hat{I}(r_0) \leq \hat{I}(\xi)$. It follows that $\hat{I}(\xi) = r_0 = \inf_{m \in \mathcal{M}} \xi(m)$. Therefore, $\hat{I}$. Now, for all $\xi, \xi' \in B(\mathcal{D}_{\mathcal{M}}$,

$$\hat{I}(\xi) - \hat{I}(\xi') = \inf_{m \in \mathcal{M}} \xi(m) - \inf_{m \in \mathcal{M}} \xi'(m) \leq \inf_{m \in \mathcal{M}} (\xi(m) - \xi(m)).$$

Thus, $\hat{I}$ is a niveloid, and it is, therefore, Lipschitz continuous. It follows that it admits a unique, monotone, and continuous extension to $B(\mathcal{D}_{\mathcal{M}})$, wich, abusing notation, we also denote $\hat{I}$. Then, pick any $\xi \in B(\mathcal{D}_{\mathcal{M}})$. Since $B_0(\mathcal{D}_{\mathcal{M}})$ is dense in $B(\mathcal{D}_{\mathcal{M}})$, we can find two sequences $(\xi_n^u)_n, (\xi_n^l)_n$ such that $\xi_n^u \searrow \xi$ and $\xi_n^l \nearrow \xi$. Since $\hat{I}$ is monotone, we have that for all $n \in \mathbb{N}$, $\xi_n^l \leq \xi \leq \xi_n^u$ and, therefore,

$$\hat{I}(\xi_n^l) = \inf_{m\in\mathcal{M}} \xi_n^l(m) \leq \inf_{m\in\mathcal{M}} \xi(m) \leq \inf_{m\in\mathcal{M}} \xi_n^u(m) = \hat{I}(\xi_n^u).$$

Since $\hat{I}$ is continuous, passing to the limit, we obtain that $\hat{I}(\xi) = \inf_{m\in\mathcal{M}} \xi$. Therefore, we have that for all $\varphi \in B_0(\mathcal{G})$,

$$\begin{aligned}
\hat{I}(I(\varphi,\cdot)) &= \inf_{m\in\mathcal{M}} \min_{p\in\Delta(\mathcal{G})} \int_\Omega \varphi dp + c(p,m) \\
&= \inf_{p\in\Delta(\mathcal{G})} \inf_{m\in\mathcal{M}} \int_\Omega \varphi dp + c(p,m) \\
&= \inf_{p\in\Delta(\mathcal{G})} \int_\Omega \varphi dp + \inf_{m\in\mathcal{M}} c(p,m).
\end{aligned}$$

Suppose that in addition $\mathcal{M}$ is closed and $\succsim$ satisfies. Fix any $\varphi \in B(\mathcal{G})$. Take any $r \in \mathbb{R}$ and pick $f \in \mathcal{F}$ and $x_r \in X$ such that $\varphi = u(f)$ and $r = u(x_r)$. Then:

$$\begin{aligned}
\{m \in \mathcal{M} : I(\varphi, m) \leq r\} &= \{m \in \mathcal{M} : I^m(u(f)) \leq u(x_r)\} \\
&= \{m \in \mathcal{M} : x_r \succsim^m f\}
\end{aligned}$$

and the latter is closed by axiom. Therefore, $m \mapsto I(\varphi, m)$ is lower semicontinuous. Therefore, the functional $\tilde{I}_\varphi : \Delta(\mathcal{G} \times \mathcal{M} \to \mathbb{R}$ defined as $\tilde{I}_\varphi(p,m) \coloneqq I(\varphi,m) - \int_\Omega \varphi dp$ is lower semicontinuous in (p, m). Then, since

$$c(p,m) = \sup_{\varphi\in B_0(\mathcal{G})} \left\{ I(\varphi,m) - \int_\Omega \varphi dp \right\} = \sup_{\varphi\in B_0(\mathcal{G})} \tilde{I}_\varphi(p,m)$$

for all $(p, m) \in \Delta \times \mathcal{M}$ and by the theorem of the maximum (see Aliprantis and Border (2007), Lemma 17.29), we can conclude that c is lower semicontinuous in

(p, m). Then applyin Aliprantis and Border (2007), Lemma 17.30 twice, we obtain that $\inf_{m\in\mathcal{M}} c(\cdot, m) = \min_{m\in\mathcal{M}} c(\cdot, m)$ is lower semicontinuous and, therefore,

$$\begin{aligned}\hat{I}(I(\varphi, \cdot)) &= \inf_{p\in\Delta(\mathcal{G})} \int_\Omega \varphi dp + \inf_{m\in\mathcal{M}} c(p, m) \\ &= \min_{p\in\Delta(\mathcal{G})} \int_\Omega \varphi dp + \min_{m\in\mathcal{M}} c(p, m).\end{aligned}$$

Since $\varphi \in B_0(\mathcal{G})$ was arbitrarily chosen, we conclude that this holds everywhere on $B_0(\mathcal{G})$. Therefore, for all $f, g \in \mathcal{F}$,

$$\begin{aligned}f \succsim g &\iff I(u(f)) \geq I(u(g)) \\ &\iff \hat{I}(I(u(f), \cdot)) \geq \hat{I}(I(u(g), \cdot)) \\ &\iff \min_{p\in\Delta(\mathcal{G})} \int_\Omega u(f)dp + \min_{m\in\mathcal{M}} c(p, m) \geq \min_{p\in\Delta(\mathcal{G})} \int_\Omega u(g)dp + \min_{m\in\mathcal{M}} c(p, m)\end{aligned}$$

as we wanted to show. ■