MISSPECIFICATION AVERSE PREFERENCES*

Alfonso Maselli†

Job Market Paper

November 29, 2025

Click here for the latest version.

Abstract. We study a decision maker who approaches a decision problem under uncertainty by formulating a set of plausible probabilistic models of the environment, while being aware that these models are only stylized and incomplete approximations. The decision maker faces two layers of uncertainty. Not only is she uncertain about which model in this set has the best fit (ambiguity), but she is also concerned that the best-fit model itself might be a poor description of the environment (model misspecification). We develop an axiomatic foundation for preferences that capture concerns about these two layers of uncertainty and allow us to compare individuals' degrees of aversion to model misspecification and to ambiguity independently of each other. In other words, these two conceptually distinct behavioral phenomena are captured by independent parameters in our representation and imply different choice patterns.

*I am particularly indebted to David Dillenberger, Kevin He, and George Mailath for their valuable guidance. I am also grateful to Aislinn Bohren, Simone Cerreia-Vioglio, Roberto Corrao, Tommaso Denti, Luigi Falasconi, Jesus Fernandez-Villaverde, Peter Klibanoff, Giacomo Lanzani, Fabio Maccheroni, Julien Manili, Massimo Marinacci, Sujoy Mukerji, Guillermo Ordonez, Andrew Postlewaite, Alex Sawyer, Ana Sofia Teles, Juuso Toikka, Gerelt Tserenjigmid, Mu Zhang, and seminar participants at UPenn and RUD 2025 for helpful comments and suggestions.
†Department of Economics, University of Pennsylvania. Email: masellia@sas.upenn.edu.

1. Introduction

Economic agents often rely on simplified and stylized descriptions of the complex environments they face to guide their decisions. This suggests that model misspecification is a pervasive phenomenon affecting many decision problems. For example, a policymaker may have an incorrect description of how the economy would respond to a fiscal or monetary stimulus, or a firm's marketing department may misjudge 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 decision making and strategic interaction (see Section 1.1 for a comprehensive literature review). A common assumption in this literature is that once agents have settled on a specific statistical model of the environment, they disregard the possibility that it is misspecified and act in a fully Bayesian fashion, evaluating alternatives by their expected utility with respect to that model. However, sufficiently sophisticated agents should realize that their model is only a simplified approximation of reality. As suggested by Hansen and Sargent (2001), an economic agent who is concerned about acting on the basis of an incorrectly formulated model should make decisions that are *robust*, that is, policies that perform reasonably well across all models that are close enough to a reference model. Following this idea, the first axiomatic treatment of decision criteria featuring misspecification aversion has been proposed by Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci (2025). Moreover, Lanzani (2025) axiomatizes a special case of preferences that are robust to misspecification and uses it to study learning under endogenous misspecification concerns.

In this paper, we provide an axiomatic foundation for a general class of preferences that capture concerns about model misspecification. Our primary contribution is to offer a clear and systematic approach to disentangle aversion to model misspecification from the more commonly studied attitudes toward ambiguity, an aspect that has not been explicitly addressed in previous work. We adopt a version of the Anscombe-Aumann framework in which uncertainty is represented by a set of states of the world Ω, and the decision maker (henceforth, DM) chooses an act f that maps states of the world to outcomes. The DM does not know the true *data-generating process (DGP)* governing the environment, but she has statistical information in the form of a set $\mathcal{M}$ of probability distributions over states of the world. Each *model* $m \in \mathcal{M}$ can be interpreted as an alternative plausible hypothesis about the DGP. Aware that models are only imperfect and stylized descriptions of the real environment, the DM may be concerned that no hypothesized model in $\mathcal{M}$ is an accurate approximation of the

DGP; equivalently, that the true DGP is not contained in $\mathcal{M}$. She also has at her disposal a best-fit map, which, for each state realization, identifies the model that best approximates the true DGP. This perspective immediately reveals two distinct channels through which uncertainty enters the decision problem:

- *ambiguity:* the agent lacks information to formulate and commit to a unique prior over the set of probabilistic models $\mathcal{M}$;
- *model misspecification:* the agent worries that the entire modeling framework is too narrow and that none of the candidate probabilistic models in $\mathcal{M}$ is a good description of the environment.

To illustrate, consider an investor who must decide how much of her wealth to allocate to a stock with an uncertain return versus a safe government bond. To solve this portfolio investment problem, the investor postulates that the stock return is normally distributed with unknown mean and variance, not because she literally believes that this assumption is correct, but because it makes the problem tractable and easier to analyze. In this scenario, the investor faces ambiguity if she is unable to pin down a unique prior over the mean and variance parameters of the hypothesized normal stock return distribution. She also faces model misspecification if she worries that the true stock return distribution might not be normal after all, but instead has thicker tails or some degree of skewness.

As the illustration highlights, ambiguity and model misspecification are conceptually distinct phenomena. The former arises from uncertainty over which probabilistic model in the hypothesized set $\mathcal{M}$ is the best description of the environment, whereas the latter pertains to the concern that the correct DGP may lie outside the scope of the set of hypothesized probabilistic models. The approach developed in this paper enables a meaningful disentangling of concerns about model misspecification from attitudes toward ambiguity, without imposing restrictive assumptions about the latter. This differs from much of the existing decision-theoretic literature, which, in order to study misspecification, has typically shut down rich ambiguity attitudes by assuming that the DM is either infinitely ambiguity averse or ambiguity neutral. By contrast, we are able to accommodate misspecification aversion while allowing complete flexibility in the DM's ambiguity attitudes (aversion, neutrality, or loving). This makes clear that misspecification aversion and ambiguity attitudes are distinct behavioral phenomena that can, in principle, lead to different patterns of choice.

Specifically, in our framework, the DM evaluates each uncertain alternative according to the following two-step procedure. First, if the DM were given sufficient information to determine that a distribution $m \in \mathcal{M}$ is the best-fit model, she would evaluate an act

f according to a quasiconcave misspecification certainty equivalent $V^m(f) = I(u(f), m)$ of the utility act induced by f, given the utility function over outcomes u. We interpret this misspecification certainty equivalent as the certain utility level that the DM would be willing to accept in order to eliminate the uncertainty stemming from the possibility of model misspecification. The quasiconcavity property of $I(\cdot, m)$ captures the DM's preference for hedging against this uncertainty and thus her concern about model misspecification. In a baseline specification of our preferences, this misspecification certainty equivalent takes the form of a misspecification robust criterion:

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

where $c(\cdot, m)$ is an index of misspecification aversion. That is, even conditional on having sufficient information to infer that m is the best-fit model, the DM does not fully trust it because of misspecification concerns. In evaluating an act f, she also takes into account distributions p outside $\mathcal{M}$ that are not too distant, in a probabilistic sense, from m. The index $c(\cdot, m)$ acts as a penalty term on this distance and captures the DM's confidence in the best-fit model m. An important special case is $c(\cdot, m) = \lambda R(\cdot || m)$, where R is the relative entropy and $\lambda > 0$ is a parameter of misspecification aversion. A lower value of λ corresponds to a higher degree of misspecification aversion, in which case the DM exhibits a preference for acts that perform robustly well across a larger set of models around m. In the limit case in which λ approaches infinity, the DM becomes misspecification neutral and evaluates acts according to their expected utility under m.

Second, because she is also uncertain about the identity of the best-fit model, the DM aggregates the misspecification robust evaluations:

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

where $\hat{I} : \mathbb{R}^{\mathcal{M}} \to \mathbb{R}$ is an aggregator capturing the DM's attitudes toward ambiguity about which model is the best-fit one. We show that $\hat{I}$ can be interpreted as a certainty equivalent of the uncertain (because of ambiguity) profile of misspecification robust evaluations; that is, the certain utility level that the DM would be willing to accept in order to eliminate ambiguity about which of the hypothesized probabilistic models has the best fit.

We illustrate how our framework distinguishes concerns about misspecification from attitudes toward ambiguity. We show that two agents can be ranked in terms of their degree of misspecification aversion by comparing only their misspecification indices c

(without imposing any mutual restrictions on their aggregators $\hat{I}$) and, symmetrically, that they can be ranked in terms of their attitudes toward ambiguity by comparing only their aggregators $\hat{I}$ (without imposing any mutual restrictions on their misspecification aversion indices). Specifically, DM1 is more misspecification averse than DM2 if and only if $c_1(\cdot, m) \leq c_2(\cdot, m)$ for all hypothesized models $m \in \mathcal{M}$. In other words, DM1 assigns lower penalties to deviations from the best-fit model and thus effectively considers a larger set of nearby alternative distributions than DM2 when evaluating acts. Conversely, DM1 is more ambiguity averse than DM2 if and only if $\hat{I}_1 \leq \hat{I}_2$. That is, DM1 is more ambiguity averse if she is willing to accept lower certainty equivalents than DM2 in order to eliminate ambiguity about the identity of the best-fit model.

We provide explicit axiomatizations of two main functional forms of the aggregator $\hat{I}$, along with several others. First, we show that if the DM confronts the uncertainty about the identity of the best-fit model in line with the expected utility tenets, she aggregates the misspecification robust evaluations in a Bayesian fashion:

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

where μ is the DM's subjective prior over the set of models $\mathcal{M}$ and ϕ captures the DM's attitudes toward ambiguity. This criterion was first proposed by Cerreia-Vioglio et al. (2025) without providing an explicit axiomatization. Moreover, when the DM is neutral toward ambiguity and exhibits a uniform concern about misspecification, this criterion reduces to the average robust control representation axiomatized by Lanzani (2025).[1] If, instead, the DM is misspecification neutral, that is, when $c(\cdot, m)$ assigns an infinite penalty to any probabilistic model different from m itself, the criterion in (3) collapses to the well-known smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005). Second, we show that if the DM 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 cautious criterion introduced by Cerreia-Vioglio et al. (2025):

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

In addition to some basic conditions, our decision criterion rests on three central axioms. The first, *coherence*, captures the idea that the DM's subjective preferences are coherent with the statistical framework embodied by the set of probabilistic models

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

and the best-fit map. Given this axiom, the DM can form complete rankings of acts conditional on the event that any given model m has the best fit. The second axiom, *misspecification aversion*, requires that even after conditioning on the event that a given $m \in \mathcal{M}$ is the best-fit model, the DM's preferences need not satisfy full-fledged independence and may display instances of uncertainty aversion. Intuitively, suppose the DM had sufficient information to determine that m is the best approximation in $\mathcal{M}$. If she were completely certain that the true DGP is included in $\mathcal{M}$, she would conclude as a matter of fact that m *is* the correct description of the environment and evaluate uncertain alternatives according to their expected utility under m. The fact that, even after m is revealed to be the best-fit model, the DM's preferences might still exhibit violations of independence and a preference for hedging against the residual uncertainty reflects her mistrust of the best-fit model m and, hence, her concern that the set of hypothesized models is misspecified. The third axiom, *consistency*, requires that the DM prefers an act f to an act g whenever f is preferred to g conditional on each $m \in \mathcal{M}$. This axiom connects the subjective preferences of the DM to the statistical information encoded in the set of models $\mathcal{M}$. It captures the idea that, even while aware of the possibility of misspecification, the DM still places substantive trust in the set of models. If they provide a unanimous ranking of two alternatives, after taking misspecification concerns into account, the DM's preferences conform to that ranking.

In the final section, we illustrate that misspecification aversion can have qualitatively different implications from ambiguity aversion by revisiting the monopoly pricing example discussed in Ball and Kattwinkel (2024). Assuming that the monopolist's preferences are represented by the criterion (4), we show that the profit guarantee of the optimal posted price is robust to perturbations of the monopolist's posited set of distributions over buyers' valuations, provided that the monopolist exhibits some degree of concern for misspecification. This stands in contrast to Ball and Kattwinkel (2024)'s finding that the optimal posted price for an infinitely ambiguity averse but misspecification unconcerned monopolist displays a sharp discontinuity in its profit guarantee, which collapses to zero under any arbitrarily small perturbation of the posited set of distributions.

The rest of the paper is structured as follows. Section 1.1 discusses the relevant literature. Section 2 presents the decision framework and introduces the set of hypothesized models and the best-fit map. Section 3 introduces and discusses the axioms characterizing the misspecification averse preferences. Section 4 states and discusses the general representation results. Section 5 characterizes various special cases of the aggregator.

Section 6 illustrates the implications of misspecification aversion in a monopoly pricing example. Section 7 concludes. All proofs are collected in the Appendix.

1.1. **Related Literature.**

Decision Criteria Incorporating Misspecification Concerns.

Following the seminal robust control model of Hansen and Sargent (2001), the first paper to axiomatize preferences that exhibit aversion to model misspecification is Cerreia-Vioglio et al. (2025). They axiomatize the criterion (4) in a two-preference setup à la Gilboa, Maccheroni, Marinacci, and Schmeidler (2010), where the DM has a family of preferences indexed by varying sets of posited models. For each posited set, the DM has an objectively rational preference that satisfies weak certainty independence but is incomplete, and a subjectively rational preference that is complete but satisfies independence only on constant acts. These two preferences are linked by two axioms introduced in Gilboa et al. (2010). The first, consistency, requires that the subjectively rational preference be a completion of the objectively rational one. The second is caution: if the objectively rational preference is not confident enough to rank an uncertain act above a deterministic one, then the deterministic act is chosen by the subjectively rational preference (when in doubt, go with the certain alternative). Moreover, the preferences are informed by the set of probabilistic models via coherence requirements similar to those used in this paper. They also examine more general aggregators of the misspecification robust evaluations and discuss the Bayesian version of the aggregator (3).

Lanzani (2025) adopts the view that states of the world can be decomposed into a payoff-relevant outcome and a component that pins down the distribution over these outcomes. He assumes that the DM has variational preferences and obtains a special case of (3), called the average robust control criterion, by requiring that preferences on bets over models satisfy the sure-thing principle and uncertainty neutrality (thereby making ϕ affine). Moreover, he proposes axioms that characterize the asymptotic behavior of the index of misspecification concern as the DM's preferences evolve in response to the arrival of new information.

We show (Theorems 4 and 5) that the cautious criterion of Cerreia-Vioglio et al. (2025) and the average robust control criterion of Lanzani (2025) fall within the general class of misspecification averse preferences studied in this paper. In particular, Cerreia-Vioglio et al. (2025)'s and Lanzani (2025)'s criteria represent two opposite ends of the spectrum: the cautious criterion displays an extreme form of ambiguity aversion, while the average robust control criterion is neutral toward ambiguity. A contribution of our paper is to allow for more flexible attitudes toward ambiguity while disentangling them from the degree of misspecification aversion. This is reflected in the fact that the representation parameters capturing model misspecification aversion (the index c) and ambiguity attitudes (the aggregator $\hat{I}$) are independent of each other.

Finally, a recent paper by Bonaglia and Dedola (2025) proposes a methodology to identify, endogenously from the preferences of a misspecification-concerned DM, the set of probabilistic models $\mathcal{M}$.

Preferences and Statistical Information.

This paper is also related to the literature connecting statistical information to choice behavior (see, for example, Amarante (2009), Al-Najjar and De Castro (2014), Epstein and Seo (2010), and Klibanoff, Mukerji, and Seo (2014)). Within this class, the closest paper to ours is Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2013). Building on their setup, we incorporate misspecification aversion into the preferences of a DM who uses exogenous statistical information to inform her choices. In their case, since the DM does not care about model misspecification, conditional on any model m she evaluates acts according to their expected utility under m. Therefore, their consistency axiom requires that whenever two acts are unanimously ranked according to their expected utility across all models m, the DM's preferences respect that ranking:

$$\forall m \in \mathcal{M},\ \mathbb{E}_m[u(f)] \geq \mathbb{E}_m[u(g)] \implies f \succsim g.$$

They show that this implies their representation depends only on the profile of expected utility evaluations $(\mathbb{E}_m[u(f)])_{m \in \mathcal{M}}$, so that preferences are represented via a monotone aggregation of this profile. In our case, however, even conditioning on the information sufficient to pin down a unique best-fit model $m \in \mathcal{M}$, the DM, out of misspecification concerns, only trusts m to be the best approximation to the DGP, but not necessarily the correct DGP itself. As a result, her preferences conditional on m need not be expected utility and may still display a preference for robustness across models in a neighborhood of m. In Theorem 1, we show that the representation of our class of misspecification averse preferences depends only on the profile of misspecification averse certainty equivalents $(I(u(f), m))_{m \in \mathcal{M}}$, so that it can be expressed as an ambiguity certainty equivalent aggregator $\hat{I}$ of the map $m \mapsto I(u(f), m)$. The fact that this map is no longer linear in the models $m \in \mathcal{M}$ is one of the main technical difficulties we address in this paper.[2] Moreover, we show that, also in our case, axioms on preferences

[2]In this respect, this paper is also related to Mu, Pomatto, Strack, and Tamuz (2024). In a different context, they show that monotone additive statistics can be represented as averages of CARA certainty equivalents.

over acts can be translated into properties of the ambiguity certainty equivalent $\hat{I}$ without having to resort to second-order acts. This allows us to readily axiomatize tractable and more easily interpretable functional forms of the aggregator $\hat{I}$.

This paper is also related to the recent axiomatization of identifiable smooth ambiguity preferences by Denti and Pomatto (2022). Without positing an exogenous set of probabilistic models and abstracting from misspecification concerns, they find conditions under which preferences are represented by the smooth ambiguity criterion, with beliefs that are identifiable; that is, they are completely orthogonal for some kernel. By contrast, we start from a DM endowed with an exogenously given set of models and a best-fit map, connect the DM's preferences to this statistical structure through coherence and consistency, and allow the DM to exhibit aversion to misspecification.

Learning with Misspecified Models.

Starting with Esponda and Pouzo (2016), many papers have examined the asymptotic behavior of actions and beliefs when agents make repeated decisions in a stochastic environment that they may understand only partially or incorrectly (see, for instance, Frick, Iijima, and Ishii, 2022; Fudenberg, Lanzani, and Strack, 2021). In all these models, agents are expected utility maximizers and do not explicitly worry about misspecification. A key result is that misspecification is asymptotically persistent and thus matters in shaping agents' behavior and beliefs, even when they collect many observations generated by the true data-generating process. A different strand of the literature allows agents to recognize that their model is misspecified and switch to a competing alternative (see, for example, Ba (2021), Fudenberg and Lanzani (2023), and He and Libgober (2021)). The main difference relative to the misspecification averse preferences axiomatized in this paper is that, once agents have selected one of the competing models on the basis of a statistical fitness test, they act in a fully Bayesian fashion.

2. Decision Framework

First, we outline the decision environment faced by the DM. Uncertainty is described by a state space Ω endowed with a countably generated sigma-algebra $\mathcal{G}$. Let X be the space of *consequences*, a convex subset of a topological vector space. The DM must choose *acts*, that is, functions $f : \Omega \to X$ mapping states to consequences that are measurable with respect to $\mathcal{G}$. An act f is bounded if there exists a finite subset K of X such that $f(\omega)$ is in the convex hull of K for all $\omega \in \Omega$. We denote by $\mathcal{F}$ the set of all bounded acts and by $\mathcal{F}_0$ the subset of acts that are simple. Abusing notation, we denote by $x \in X$ the constant act yielding consequence x in each state $\omega \in \Omega$. For each $f, g \in \mathcal{F}$ and $\alpha \in [0,1]$, the convex combination $\alpha f + (1-\alpha)g$ is the act given by:

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

for all $\omega \in \Omega$. Given any $E \in \mathcal{G}$ and acts $f, g \in \mathcal{F}$, let fEg be the act taking the value $f(\omega)$ if $\omega \in E$ and the value $g(\omega)$ if $\omega \in \Omega \setminus E$. If $\mathcal{E}$ is a sub-sigma-algebra of $\mathcal{G}$, denote by $\mathcal{F}(\mathcal{E})$ the subset of acts in $\mathcal{F}$ that are measurable with respect to $\mathcal{E}$.

Let $\succsim$ be a preference relation over $\mathcal{F}$. We denote by $\succ$ and $\sim$, respectively, the asymmetric and symmetric parts of $\succsim$. An event $E \in \mathcal{G}$ is *null* if $fEh \sim gEh$ for all acts $f, g, h \in \mathcal{F}$ and *nonnull* otherwise.

2.1. **Probabilistic Models and Best-Fit Map.** Let $\Delta \coloneqq \Delta(\Omega, \mathcal{G})$ denote the space of countably additive probability measures on $(\Omega, \mathcal{G})$. Endow Δ with the natural sigma-algebra $\mathcal{D}$ generated by the family of evaluation maps and any subset of Δ, with its relative sigma-algebra.[3]

We assume that the DM is equipped with a set $\mathcal{M} \subseteq \Delta$ of probability distributions over states of the world that, given some external information, she believes are plausible descriptions of the uncertain environment she faces. In line with the classical setup of Wald (1950), each model $m \in \mathcal{M}$ can be interpreted as an alternative hypothesis about the DGP, grounded in substantive considerations such as scientific theories and empirical evidence. The models in $\mathcal{M}$ are sometimes referred to in the literature (see, for example, Hansen and Sargent (2022) and Cerreia-Vioglio et al. (2025)) as *structured* to emphasize their special status in the eyes of the DM relative to other distributions outside $\mathcal{M}$. Motivated by the view, dating back to Box (1976, 1979) and Cox (1995), that models are only approximations, we do not assume that the set of models contains the DGP, that is, the true probability law governing state uncertainty. Moreover, we allow for the possibility that the DM is aware of this and recognizes that her set of models may be *misspecified*.

Our aim in this paper is to distinguish between ambiguity about which probabilistic model best approximates the DGP and concern about misspecification, namely the possibility that no hypothesized model is an accurate approximation of the DGP. Uncertainty about models is usually attributed to a "lack of information" that prevents the DM from selecting a single best model. Following Cerreia-Vioglio et al. (2013),[4] we formalize this lack of information using the notion of sufficient statistics in the sense

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

[4]See also Amarante (2009), Al-Najjar and De Castro (2014), Epstein and Seo (2010), and Klibanoff et al. (2014) for related approaches and Denti and Pomatto (2022) for a discussion of this condition.

of Dynkin (1978). Specifically, we assume that the measurable space of states of the world $(\Omega, \mathcal{G})$ and the set of probabilistic models $\mathcal{M}$ admit a (measurable) *best-fit map* $\mathfrak{q} : \Omega \to \mathcal{M}$ such that

$$m(\{\omega \in \Omega : \mathfrak{q}(\omega) = m\}) = 1 \text{ for all } m \in \mathcal{M}. \tag{5}$$ [5]

We interpret the event $E^m = \{\omega : \mathfrak{q}(\omega) = m\}$ in which model m is selected by the map $\mathfrak{q}$ as the event that m is the probabilistic model in $\mathcal{M}$ that most closely resembles the true distribution over the states of the world. The sigma-algebra $\mathcal{A}$ generated by the best-fit map $\mathfrak{q}$ can thus be interpreted as representing *sufficient* information to determine which model in the family $\mathcal{M}$ is the closest approximation to the true DGP, that is, the one with the best fit. Condition (5) is the well-known assumption that the DM's family of probabilistic models is *point identified* whenever it is correctly specified. In particular, this condition only requires that if a model $m \in \mathcal{M}$ coincides with the true DGP, then the best-fit map selects it with probability one.

We interpret this framework as follows. Suppose that a well-specified description of the environment is given by a set of models $\mathcal{P}$ and a map $p : \omega \mapsto p^\omega \in \Delta$ such that $P(\{\omega \in \Omega : p^\omega = P\}) = 1$ for all $P \in \mathcal{P}$. The interpretation is that the realization ω also determines the true DGP p^ω. Moreover, the statement that $p^\omega = P$ with probability one under P ensures that the description of the environment is internally consistent; that is, whenever P is the true DGP, it is selected with probability one by the map p.[6] However, the DM posits a misspecified set of models $\mathcal{M} \subseteq \mathcal{P}$, which does not necessarily include every model in $\mathcal{P}$. Suppose that the DM attempts to estimate the best-fit model within the hypothesized set $\mathcal{M}$. Because the DM has posited a misspecified set of models, the insight from Berk (1966) suggests that she would asymptotically select the model in $\mathcal{M}$ that is closest to P; that is, the (assumed unique) minimizer $q^*(P) \in \mathcal{M}$ solving $\min_{m \in \mathcal{M}} D(P||m)$, where $D(\cdot||\cdot)$ is an appropriate measure of fit. Define the function $\mathfrak{q}(\omega) \coloneqq q^*(p^\omega)$. Then the set E^m represents the event that m has the best fit, and the sigma-algebra $\mathcal{A}$ generated by $\mathfrak{q}$ encodes the information needed to determine which hypothesized model has the best fit. Moreover, for all $m \in \mathcal{M}$, we have that $m\left(\{\omega : \mathfrak{q}(\omega) = m\}\right) = 1$. It is in this sense that we interpret $\mathfrak{q}$ as a *best-fit*

[5] The requirement that each model $m \in \mathcal{M}$ is selected by the best-fit map with probability one under m is equivalent to the notion of sufficient statistics introduced by Dynkin (1978) and is related to the strong law of large numbers. Mathematically, this property is known as complete orthogonality of a set of probabilistic models $\mathcal{M}$. See, for example, Mauldin, Preiss, and Weizsacker (1983) and Weis (1984).

[6] We can see an 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.

map and the information in $\mathcal{A}$ as sufficient information to determine the best approximation to the DGP among those in $\mathcal{M}$. That is, if the DM were able to observe ω, she would infer that the model $m^\omega = \mathfrak{q}(\omega)$ is the model that most closely resembles the true DGP.

Example 1: Recall the portfolio investment example mentioned in the introduction, and consider an investor who needs to decide at the initial period how much to invest in a stock versus a safe perpetuity paying a fixed amount in every future period. At each period t, a stock return $s_t \in \mathbb{R}$ is realized. From the investor's perspective at the initial date, a state of the world is an infinite sequence of stock return realizations, and the state space is given by the sequence space $\Omega = \mathbb{R}^{\mathbb{N}}$. To solve the portfolio problem, the investor postulates that stock returns are i.i.d. normally distributed over time, with unknown mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in \mathbb{R}_{++}$. She makes this assumption not because she believes it is inherently correct, but because it provides a tractable, easy-to-use model for solving the investment problem. In this scenario, the set of probabilistic models posited by the investor is

$$\mathcal{M} = \{m_{\mu,\sigma^2} = \times_{n=1}^{\infty} \Phi_{\mu,\sigma^2} : \Phi_{\mu,\sigma^2} \sim N(\mu, \sigma^2),\ \mu \in \mathbb{R},\ \sigma^2 > 0\}.$$

Suppose, however, that the correct description of the environment is that the sequence of stock returns is i.i.d. with marginal distribution given by a generalized normal distribution p_θ, parametrized by $\theta = (\theta_1, \theta_2, \theta_3) \in \mathbb{R} \times \mathbb{R}^2_{++}$, which respectively capture the mean, the variance, and the tail thickness of the distribution. The investor attempts to parse the uncertainty regarding the mean and variance of the (misspecified) normal family of models she posited by considering the asymptotic behavior of the maximum likelihood estimators based on the observations of stock returns accumulated over time. She then constructs her best-fit map $\mathfrak{q} : \Omega \to \mathbb{R} \times \mathbb{R}_{++}$ as the limit of these estimates. By the results in White (1982), we know that, with probability one under the true i.i.d. generalized normal DGP with parameter θ, this limit equals the unique minimizer $(\hat{\mu}, \hat{\sigma}^2)$ solving $\min_{\mu,\sigma^2} R(p_\theta, \Phi_{\mu,\sigma^2})$, which in turn matches exactly the mean and variance of the true DGP: $(\hat{\mu}, \hat{\sigma}^2) = (\theta_1, \theta_2)$.[7] We can thus interpret the i.i.d. normal model with mean θ_1 and variance θ_2 as the best approximation to the true DGP. This would be true whenever the generalized normal DGP has the first two parameter components equal to (θ_1, θ_2), regardless of the value of the third parameter θ_3, which controls tail thickness (including the case in which the DGP is itself i.i.d. normal, corresponding to the case $\theta_3 = 2$). Thus, we can interpret the event $\{\omega : \mathfrak{q}(\omega) = (\theta_1, \theta_2)\}$

[7]Here R is the relative entropy; that is, $R(q||p) = \int_\Omega \ln \frac{dq}{dp} dq$ if $q \ll p$ and equal to $+\infty$ otherwise.

as the event that the i.i.d. normal model with mean θ_1 and variance θ_2 is the model in the hypothesized set that has the best fit and $\mathfrak{q}$ as a best-fit map.

3. Misspecification Averse Preferences Axioms

In this section, we fix a measurable state space $(\Omega, \mathcal{G})$ and a set of probabilistic models $\mathcal{M}$ that admit a best-fit map $\mathfrak{q}$ satisfying the properties outlined in the previous section.

3.1. **Basic Conditions.**

We assume that the DM's preferences satisfy some basic conditions, which are discussed, for example, in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011).

Axiom 1 (Basic Conditions):

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

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

(iii) Mixture Continuity. *If* $f, g, h \in \mathcal{F}$, *the sets* $\{\alpha \in [0,1] : \alpha f + (1-\alpha) g \succsim h\}$ *and* $\{\alpha \in [0,1] : h \succsim \alpha f + (1-\alpha) g\}$ *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) 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 guarantee that the preferences are a continuous and monotone weak order and satisfy independence when restricted to constant acts. By the theorem of Herstein and Milnor (1953), these preferences are represented on X by an affine utility function u. If we interpret the mixture space X as the set of simple lotteries over outcomes, these axioms imply that the DM evaluates lotteries, that is, constant acts that do not involve uncertainty about the state of the world, according to their objective expected utility. The last requirement is mostly for technical convenience and guarantees that the utility over consequences u is unbounded above and below.

3.2. **Coherence.**

For each model $m \in \mathcal{M}$, recall that we defined $E^m := \mathfrak{q}^{-1}(m) \in \mathcal{A}$ as the set of states of the world for which the best-fit map selects m as the best approximation to the DGP. The following axiom captures the idea that the DM's preferences are coherent with the statistical framework embodied by the set of models and the best-fit map.

Axiom 2 (Coherence):

(i) For all models $m \in \mathcal{M}$*, for all* $f, g, h, h' \in \mathcal{F}$*,*

$$fE^m h \succsim gE^m h \iff fE^m h' \succsim gE^m h',$$

and $x \succsim yE^m x$ *if and only if* $x \succsim y$ *for all* $x, y \in X$*.*

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

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

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

Point *(i)* has two components. First, it requires that each event E^m satisfies Savage's *sure-thing principle*. This means that for each $m \in \mathcal{M}$, the DM can identify the event in which model m is the best approximation to the DGP and can evaluate acts conditional on this event. In particular, this allows us to define, in an unambiguous way, a conditional preference relation $\succsim^m$ given that a model $m \in \mathcal{M}$ is the best approximation to the true DGP, by requiring that, for all $f, g \in \mathcal{F}$,

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

Second, it requires that preferences conditional on the event that a model m has the best fit rank constant acts in the same way as the original preference $\succsim$. In other words, the DM's preferences over consequences are stable across different event realizations.

Point *(ii)* ensures that the DM's preferences incorporate the information provided by the best-fit map and are coherent with the selected best-fit model. If two acts are equal with probability one under a model $m \in \mathcal{M}$, then the fact that the DM pays special attention to this model when it is the best-fit one suggests that she will be indifferent between them, conditional on the event that m is indeed the best approximation.

Point *(iii)* is a technical continuity requirement on how model-conditional preferences vary over the set of hypothesized models.

Taken together, these three coherence conditions imply that each model $m \in \mathcal{M}$ induces a well-defined and non-trivial conditional preference $\succsim^m$, which ranks as indifferent any acts that are equal with probability one under m, and that the model-conditional preferences $\succsim^m$ vary continuously with the model m.

3.3. **Misspecification Aversion and Consistency.** We now state the axiom characterizing the DM's aversion to misspecification. It is a model-conditional version of the uncertainty aversion axiom introduced by Schmeidler (1989). In particular, preferences conditional on the event that a model $m \in \mathcal{M}$ is the best-fit need not satisfy full-fledged independence, but they still display a preference for hedging against the uncertainty generated by model misspecification.

Axiom 3 (Misspecification Aversion): *For all models $m \in \mathcal{M}$, $f, g \in \mathcal{F}$, and $\alpha \in (0,1)$,*

$$f \sim gE^m f \implies \alpha f + (1-\alpha) gE^m f \succsim f \ .$$

We interpret Axiom 3 as capturing the idea that the DM is aware that her set of models may be 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 that the DM were able to observe sufficient information to determine that a model m is the best-fit among all those in $\mathcal{M}$. If she were fully confident that the true DGP is included in $\mathcal{M}$, she would then conclude that m *is* the correct description of uncertainty over states. In that case, there would be no reason for the DM's preferences to display any uncertainty aversion; conditional on m, they should be consistent with subjective expected utility. Hence any residual uncertainty aversion can only be ascribed to the DM's concern that her set of probabilistic models is misspecified. This is precisely the content of Axiom 3: even after being told that m is the best-fit model, the DM still exhibits violations of independence and a preference for hedging against the residual uncertainty, thereby revealing her concern for model misspecification.

The next axiom plays a crucial role in linking the DM's subjective preferences to the set of models and the conditional preferences they induce.

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

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

This assumption is analogous to the consistency axiom introduced in Gilboa et al. (2010) and Cerreia-Vioglio et al. (2013) and captures the idea that the DM uses the set of hypothesized models to inform her preferences. If an act f is unanimously ranked above an act g conditional on every model $m \in \mathcal{M}$, then the DM prefers f to g.

Axioms 3 and 4 epitomize our perspective on how to identify the DM's misspecification aversion. In our approach, concerns about misspecification are evaluated model by model. Axiom 3 states that, for each model m, the DM's preferences conditional on m being the best-fit model are averse to the residual uncertainty arising from the possibility that m is not an accurate description of the environment. Having incorporated misspecification concerns for each hypothesized model, the DM then lets the misspecification averse, model-conditional rankings $\succsim^m$ inform her overall subjective preference $\succsim$ via the consistency axiom. Only at this stage, after misspecification concerns have been dealt with "pointwise" across models, does the DM express attitudes toward ambiguity about which model has the best fit.

To summarize, we define the preferences under analysis as a binary relation that satisfies all the axioms discussed so far.

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

Before turning to the representation results in the next section, note that we have not yet imposed any restriction on the DM's attitudes toward ambiguity. Our approach enables us to incorporate the DM's concerns about misspecification while leaving her attitudes toward ambiguity entirely flexible. The following definition classifies the DM's attitudes toward ambiguity in terms of her hedging behavior with respect to uncertainty about which model in $\mathcal{M}$ has the best fit.

Definition 2 (Ambiguity Attitudes): We say that the DM's preference $\succsim$ is:

- *ambiguity averse* if $f \sim g$ implies that $\alpha f + (1-\alpha) g \succsim f$ for all $f, g \in \mathcal{F}(\mathcal{A})$ and $\alpha \in (0,1)$;
- *ambiguity loving* if $f \sim g$ implies that $f \succsim \alpha f + (1-\alpha) g$ for all $f, g \in \mathcal{F}(\mathcal{A})$ and $\alpha \in (0,1)$.

Since uncertainty about which model has the best fit is represented by the events in the sigma-algebra $\mathcal{A}$, the DM's attitudes toward ambiguity are fully encoded in her preferences over acts that are $\mathcal{A}$-measurable. Thus, the DM exhibits ambiguity aversion if her preferences are uncertainty averse on $\mathcal{F}(\mathcal{A})$, that is, if she displays a preference for hedging against uncertainty about which model has the best fit. Conversely, the DM is ambiguity loving if the opposite holds, that is, if she dislikes hedging against this uncertainty.

4. Representation of Misspecification Averse Preferences

We say that a map $I : B(\mathcal{G}) \times \mathcal{M} \to \mathbb{R}$ is a *family of misspecification averse certainty equivalents* if it is lower semicontinuous in its second argument and, for all $m \in \mathcal{M}$, the function $I(\cdot, m)$ is monotone, normalized, continuous, and quasiconcave. We say that a map $\hat{I} : B(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ is an *ambiguity certainty equivalent aggregator* if it is monotone, normalized, and continuous.[8]

THEOREM 1: *Suppose that the triple* $(\Omega, \mathcal{G}, \mathcal{M})$ *admits a best-fit map* $\mathfrak{q}$*. The following statements are equivalent:*

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

(ii) there exist an affine and surjective utility function $u : X \to \mathbb{R}$*, a family of misspecification averse certainty equivalents* $I : B(\mathcal{G}) \times \mathcal{M} \to \mathbb{R}$ *and an ambiguity certainty equivalent aggregator* $\hat{I} : B(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$*, such that for all* $m \in \mathcal{M}$*,*

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

and

$$f \succsim g \iff \hat{I}\left(I(u(f), \cdot)\right) \geq \hat{I}\left(I(u(g), \cdot)\right) \tag{6}$$

for all $f, g \in \mathcal{F}$*. In addition, for all* $m \in \mathcal{M}$*,* $\varphi = \varphi'$ *a.e.* $[m] \implies I(\varphi, m) = I(\varphi', m)$ *for all* $\varphi, \varphi' \in B(\mathcal{G})$*.*

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

We can interpret the representation of the misspecification averse preferences characterized in Theorem 1 as a two-step procedure for evaluating an act. First, suppose the DM is told that a model m is the best-fit model within $\mathcal{M}$. In this case, that is, conditional on the event E^m, she evaluates an act f according to the misspecification averse certainty equivalent $I(u(f), m)$, where u is a utility function over consequences. In other words, the quantity $I(u(f), m)$ is the constant utility level the DM would be willing to accept in order to eliminate uncertainty when she is sure that m is the best-fit model. Conditional on knowing that m has the best fit, any residual uncertainty can only be attributed to possible misspecification of m. The fact that the certainty equivalent $I(\cdot, m)$ is quasiconcave, and in general is not simply the expected value of the utility act $u(f)$ under model m, therefore reflects the DM's concern about misspecification. Each map $I(\cdot, m)$ can be viewed as a non-linear expectation with respect to the best-fit model $m \in \mathcal{M}$. Although it fails to be linear, it still satisfies many of the characteristic properties of expectations, such as monotonicity, normalization, and identical evaluation of functions that are almost surely equal under m. In particular,

[8] We denote by $B(\mathcal{G})$ the set of measurable and bounded functions $\varphi : \Omega \to \mathbb{R}$. Similarly, $B(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$ denotes the set of measurable and bounded functions $\varphi : \mathcal{M} \to \mathbb{R}$. In the remainder of the paper, continuity is always understood with respect to convergence in the supnorm on the relevant space of bounded functions.

for each function φ, the random variable $\omega \mapsto I(\varphi, \mathfrak{q}(\omega))$ can be interpreted as a nonlinear "common" conditional expectation of φ given the sigma-algebra $\mathcal{A}$ for the family $\{I(\cdot, m)\}_{m\in\mathcal{M}}$.[9]

Given the representation of the model-conditional preferences, we can associate with each act $f \in \mathcal{F}$ the function $m \mapsto I(u(f), m)$, which maps each hypothesized model m to the misspecification averse certainty equivalent of act f conditional on m being the best-fit model. The axiom of consistency then implies that if $I(u(f), m) \geq I(u(g), m)$ for all $m \in \mathcal{M}$, the DM is confident that f is at least as good as g and therefore $f \succsim g$. As remarked in Cerreia-Vioglio et al. (2025), this is a manifestation of the special status of the hypothesized models in $\mathcal{M}$. When the misspecification averse certainty equivalents under each model m unanimously rank one act above another, this unanimity is sufficient for the DM to choose the former act. In general, however, the set of models will not yield a unanimous ranking for every pair of acts. In a second step, the DM therefore aggregates the misspecification averse certainty equivalents of an act f by means of the aggregator $\hat{I}$, an ambiguity certainty equivalent that captures her attitudes toward ambiguity about which model is the best-fit one. The fact that we have not taken a stance on the DM's ambiguity attitudes is reflected in the representation by the absence of any restriction on the curvature of the aggregator $\hat{I}$. The following result makes explicit how the aggregator $\hat{I}$ encodes the DM's attitudes toward ambiguity.

Corollary 1: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$ and $\succsim$ is a misspecification averse preference relation. Then:*

- *$\succsim$ is ambiguity averse if and only if $\hat{I}$ is quasiconcave;*
- *$\succsim$ is ambiguity loving if and only if $\hat{I}$ is quasiconvex.*

The representation of the misspecification averse preferences in Theorem 1 fully disentangles misspecification concerns from attitudes toward ambiguity. Misspecification aversion is entirely captured by the family of misspecification averse certainty equivalents $\{I(\cdot, m)\}_{m\in\mathcal{M}}$, independently of the shape of the ambiguity aggregator. Conversely, attitudes toward ambiguity are entirely captured by the aggregator $\hat{I}$, independently of the family of misspecification averse certainty equivalents. We now formalize this separation by means of a comparative statics exercise. To this end, we adapt the notion of comparative uncertainty aversion introduced by Ghirardato and Marinacci (2002) to the present context. Consider two decision makers, DM1 and DM2, with preference relations $\succsim_1$ and $\succsim_2$, respectively. We say that $\succsim_1$ is *more*

[9]That is, an analogue of the law of iterated expectations holds: $I\left(I(\varphi, \mathfrak{q}(\cdot)), m\right) = I(\varphi, m)$ for all $m \in \mathcal{M}$.

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 is that constant acts are unaffected by the possibility that the set of hypothesized models is misspecified, since they are non-stochastic and their evaluation does not depend on the probabilistic assessment of state uncertainty. Hence, if for every model $m \in \mathcal{M}$, DM1 is not sufficiently concerned about misspecification to choose the constant act x over the uncertain act $fE^m x$, then, a fortiori, the same must hold for the less misspecification averse DM2.

On the other hand, we say that $\succsim_1$ is *more ambiguity averse* than $\succsim_2$ if for all $f \in \mathcal{F}(\mathcal{A})$ and $x \in X$,

$$f \succsim_1 x \implies f \succsim_2 x \; . \tag{8}$$

The intuition is that the acts in $\mathcal{F}(\mathcal{A})$ are precisely those that depend only on uncertainty about which model in $\mathcal{M}$ is the best approximation and are not affected by misspecification concerns (in particular, such acts must be constant on each event E^m). Thus, if $\succsim_1$ is more ambiguity averse than $\succsim_2$, then whenever ambiguity considerations are not strong enough for DM1 to prefer the certain outcome x to the act f, whose payoff is affected solely by ambiguity about the best-fit model, they cannot be strong enough for the less ambiguity averse DM2.

We then obtain the following result, which characterizes the comparative statics of misspecification aversion and attitudes toward ambiguity in terms of the components of the misspecification averse representation.

Proposition 1: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$ and $\succsim_1$ and $\succsim_2$ are two misspecification averse preference relations. Then:*

- *(i) $\succsim_1$ is more misspecification averse than $\succsim_2$ if and only if u_2 is a positive affine transformation of u_1 and, after normalizing $u_1 = u_2$, $I_1(\cdot, m) \leq I_2(\cdot, m)$ for all models $m \in \mathcal{M}$.*
- *(ii) $\succsim_1$ is more ambiguity averse than $\succsim_2$ if and only if u_2 is a positive affine transformation of u_1 and, after normalizing $u_1 = u_2$, $\hat{I}_1 \leq \hat{I}_2$.*

The first part of the result states that $\succsim_1$ is more averse to model misspecification than $\succsim_2$ if, conditional on any given model m having the best fit, DM1 is willing to accept lower certainty equivalents than DM2 as compensation for acts that are exposed to the possibility of misspecification. The second part of the result can be read as saying that $\succsim_1$ is more averse to ambiguity than $\succsim_2$ if DM1 is willing to accept lower certainty equivalents than DM2 as compensation for uncertain bets on which model in $\mathcal{M}$ is the best approximation. Thus, Proposition 1 clarifies how representation (6) separates attitudes toward ambiguity about the identity of the best-fit model from concerns about misspecification. Attitudes toward ambiguity are captured by the aggregator $\hat{I}$, while the family of misspecification averse certainty equivalents $\{I(\cdot, m)\}_{m\in\mathcal{M}}$ serves as an index of the degree of aversion to the possibility that the set of hypothesized models is misspecified.

In the remainder of the paper, we focus on a special class of misspecification averse certainty equivalents that allows us to interpret the DM's concerns about misspecification in terms of a robust approach to model misspecification. This special case is obtained by adding more structure to the preferences conditional on each model m having the best fit. Before introducing such structure, we impose an axiom due to Arrow (1970), which guarantees that preferences are robust to small perturbations and ensures countable additivity of the subjective probabilities involved in the representation results.

Axiom 5 (Monotone Continuity): *For all* $f, g \in \mathcal{F}$ *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* $f \succ g$, *then, there exists* $n_0 \in \mathbb{N}$ *such that* $xE_{n_0}f \succ g$.

Now, define, for each probability distribution $p \in \Delta$ and each simple act $f \in \mathcal{F}_0$, the "average" of f under p by:

$$\mathbb{E}_p[f] \coloneqq \sum_{x\in X} xp\left(f^{-1}(x)\right).$$

Since f has a finite image and X is convex, $\mathbb{E}_p[f] \in X$ is the constant act that an Anscombe-Aumann expected utility maximizer with belief p over the state space Ω would regard as indifferent to f.

Axiom 6 (Variational Misspecification):

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

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

(ii) *For all* $m \in \mathcal{M}$, *if* $p(E^m) = 1$ *but* $p \neq m$, *then there exist* $f \in \mathcal{F}_0$ *and* $x \in X$ *such that* $fE^m x \succsim x$ *but* $x \succ \mathbb{E}_p[fE^m x]$.

The first part of the axiom is a model-conditional analogue of the axiom characterizing variational preferences in Maccheroni, Marinacci, and Rustichini (2006). Conditional on the event that a model $m \in \mathcal{M}$ is the best-fit one, preferences satisfy a weaker form of independence, namely weak certainty independence. Because of misspecification concerns, however, they still need not satisfy full-fledged independence.

The second part of the axiom clarifies the special status of each hypothesized model $m \in \mathcal{M}$ in the eyes of the DM relative to probabilistic models outside $\mathcal{M}$. In principle, any probability measure p that assigns probability one to E^m is statistically consistent with observing E^m. However, because such a p is not selected by the best-fit map $\mathfrak{q}$, the DM's preferences are "incoherent" with p. In particular, the DM may still prefer a (possibly) uncertain act f to a constant act x conditional on E^m, even though x is strictly preferred to the p-average of f.[10] Armed with this additional axiom, we obtain the following result.

THEOREM 2: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$. The following statements are equivalent:*

(i) $\succsim$ is a misspecification averse preference relation satisfying Axioms 5 and 6,

(ii) there exist an affine and surjective utility function $u : X \to \mathbb{R}$, a convex statistical divergence distance $c : \Delta \times \mathcal{M} \to [0, \infty]$, and a monotone continuous ambiguity certainty equivalent aggregator $\hat{I} : B(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$,[11] such that for all $f, g \in \mathcal{F}$,

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

and for all $m \in \mathcal{M}$, $\succsim^m$ is represented by

$$I(u(f), m) = \min_{p \ll m} \left\{ \int_\Omega u(f)dp + c(p, m) \right\} \quad \text{for all } f \in \mathcal{F}. \tag{9}$$

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

We can interpret the functional form (9) as embodying a robust approach to the possibility of misspecification. Suppose that the DM is sure that m is the best-fit model in her hypothesized set $\mathcal{M}$. Because she fears misspecification, she does not rely solely on m when evaluating an act f conditional on this information; instead, she forms a robust evaluation of f by also taking into account alternative probability distributions p outside $\mathcal{M}$, penalizing each such p according to its statistical distance from m. The distance $c(\cdot, m)$ thus encodes her confidence in the benchmark model m: when $c(\cdot, m)$ is

[10] This last requirement is not strictly needed for the representation result in Theorem 2. Its role is to clarify the interpretation of each m as the unique reference model for the DM conditional on the event E^m. In particular, it implies that the misspecification index $c(\cdot, m)$ in Theorem 2 is uniquely minimized at m.

[11] See Appendix A for a rigorous definition of the notion of statistical divergence distance and the property of monotone continuity of an operator.

(uniformly) lower, the penalty for deviating from m is smaller, so the DM allows a larger set of nearby models around m to influence the evaluation of f. This reflects lower trust in m or, equivalently, greater aversion to misspecification. An important and tractable specification is when the distance takes the form $c(\cdot, m) = \lambda R(\cdot||m)$ for all hypothesized models $m \in \mathcal{M}$, where $\lambda > 0$ is a parameter governing misspecification aversion and R denotes the relative entropy. In this case, misspecification concerns are proportional to the relative entropy from the best-fit model and are uniform across models in $\mathcal{M}$ (see Lanzani (2025)), with higher aversion to misspecification corresponding to a lower value of λ.

We now formalize this interpretation of the statistical distance $c(\cdot, m)$ as an index of the degree of misspecification aversion by means of a comparative statics exercise.

COROLLARY 2: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$ and $\succsim_1$ and $\succsim_2$ are two misspecification averse preference relations satisfying Axioms 5 and 6. 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 models $m \in \mathcal{M}$.*

In this paper, we treat the best-fit map as an exogenous object available to the DM. However, the next result shows that if preferences are represented by the criterion in Theorem 2, then there exists a best-fit map $\hat{\mathfrak{q}}$ with respect to which the preferences satisfy the coherence, misspecification aversion, and consistency axioms. To state it, we introduce the following definition. The aggregator $\hat{I}$ is *strongly monotone* if for all $\xi_1, \xi_2 \in B(\mathcal{M})$ such that $\xi_1 > \xi_2$[12], then $\hat{I}(\xi_1) > \hat{I}(\xi_2)$.

PROPOSITION 2: *Assume that $(\Omega, \mathcal{G})$ is a standard Borel space and that $\mathcal{M}$ is a measurable subset of Δ (not assumed to admit a best-fit map). Suppose that a preference relation $\hat{\succsim}$ is represented by the criterion characterized in Theorem 2 and that for all $\xi \in B_0(\mathcal{M})$ such that $0 \leq \xi \leq 1$ there exists $\varphi \in B_0(\mathcal{G})$ such that $0 \leq \varphi \leq 1$ and $\xi(m) = I(\varphi, m)$ for all $m \in \mathcal{M}$. Then, there exists a best-fit map $\hat{\mathfrak{q}} : \Omega \to \mathcal{M}$ such that $m\left(\hat{\mathfrak{q}}^{-1}(m)\right) = 1$ for all $m \in \mathcal{M}$ and if we define $\hat{E}^m = \hat{\mathfrak{q}}^{-1}(m)$, then*

$$f = g \quad a.e.\ [m] \implies f\hat{E}^m h \hat{\sim} g\hat{E}^m h$$

for all $f, g, h \in \mathcal{F}$. If, furthermore, $\hat{I}$ is strongly monotone, then $\hat{\succsim}$ satisfies Axioms 2, 3, and 4 given the best-fit map $\hat{\mathfrak{q}}$.

In the general representations in Theorems 1 and 2, the abstract nature of $\hat{I}$ reflects the fact that we impose no behavioral axioms on how the DM confronts uncertainty

[12]Recall that $\xi_1 > \xi_2$ if $\xi_1 \geq \xi_2$ and $\xi_1(m) > \xi_2(m)$ for some $m \in \mathcal{M}$.

about which model in $\mathcal{M}$ has the best fit. As argued above, the sigma-algebra $\mathcal{A}$ generated by the best-fit map $\mathfrak{q}$ captures the sufficient information needed to determine which model in $\mathcal{M}$ is the best approximation available to the DM. We can therefore interpret the set $\mathcal{F}(\mathcal{A})$ of acts that are measurable with respect to $\mathcal{A}$ as bets over the identity of the best-fit model. Since the aggregator $\hat{I}$ captures only the DM's attitudes toward ambiguity, it is natural to expect that its functional form depends solely on the DM's preferences over this subset of acts. Accordingly, in the next section we characterize special cases of the aggregator by imposing suitable axioms on the restriction of preferences to $\mathcal{F}(\mathcal{A})$.

5. Special Cases of the Ambiguity Aggregator

We first consider the case in which the aggregator $\hat{I}$ itself takes a variational form. To this end, we introduce the following axiom, which requires the restriction of preferences to $\mathcal{F}(\mathcal{A})$ to satisfy Maccheroni et al. (2006)'s weak certainty independence as well as Schmeidler (1989)'s uncertainty aversion.

Axiom 7 ($\mathcal{A}$-Variational):

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

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

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

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

This leads to the following result.

Theorem 3: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$. The following statements are equivalent:*

(i) $\succsim$ is a misspecification averse preference relation satisfying Axioms 5, 6, and 7,

(ii) there exist an affine and surjective utility function $u : X \to \mathbb{R}$, a convex statistical divergence distance $c : \Delta \times \mathcal{M} \to [0,\infty]$, and a grounded, lower semi-continuous and convex function $\kappa : \Delta(\mathcal{M}) \to [0,\infty]$ such that $\succsim$ is represented by:

$$\min_{\nu \in \Delta(\mathcal{M})} \left\{ \int_{\mathcal{M}} \min_{p \in \Delta(\Omega)} \{\mathbb{E}_p[u(f)] + c(p,m)\} \, d\nu(m) + \kappa(\nu) \right\} \tag{10}$$

and $\min_p \{\mathbb{E}_p[u(f)] + c(p,m)\}$ represents $\succsim^m$ for all $m \in \mathcal{M}$.

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

In representation (10), ambiguity manifests itself in the DM's inability to formulate a unique prior ν over models in $\mathcal{M}$. The convex cost $\kappa : \Delta(\mathcal{M}) \to [0, +\infty]$ serves as an index of ambiguity aversion: a lower κ corresponds to a higher degree of ambiguity aversion. Note that the functional form of the aggregator characterized in Theorem 3 also encompasses the seminal maxmin expected utility model of Gilboa and Schmeidler (1989). Indeed, if we let $\kappa(\nu) = 0$ for all ν in a compact and convex set of priors $C \subseteq \Delta(\mathcal{M})$ over the set of hypothesized models and let it be ∞ otherwise,[13] criterion (10) takes the form:

$$\min_{\nu \in C} \int_{\mathcal{M}} \min_{p \in \Delta(\Omega)} \{\mathbb{E}_p[u(f)] + c(p,m)\}\, d\nu(m) \ .$$

Because of ambiguity, the DM cannot pin down a unique prior over the set of models $\mathcal{M}$. Instead, she can only form a set of possible priors C and, out of ambiguity aversion, evaluates acts according to the worst-case expected value of the misspecification robust certainty equivalents.

Another tractable instance of criterion (10) arises when both c and κ are proportional to the relative entropy. As shown by Strzalecki (2011)'s axiomatization of the multiplier preferences introduced by Hansen and Sargent (2001), this is the case when preferences satisfy Savage's sure-thing principle.

Axiom 8 (Savage's P2. Sure-Thing Principle): *For all $E \in \mathcal{G}$ and $f, g, h, h' \in \mathcal{F}$, $fEh \succsim gEh$ if and only if $fEh' \succsim gEh'$.*

We then obtain the following corollary.

Corollary 3: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$, that $|\mathcal{M}| > 3$, and that E^m contains at least three disjoint nonnull events for all $m \in \mathcal{M}$. Then, $\succsim$ is a misspecification averse preference relation satisfying Axioms 5, 6, 7, and 8 if and only if there exist an affine and surjective utility function $u : X \to \mathbb{R}$, misspecification aversion parameters $\lambda_m \in [0, \infty]$ for all $m \in \mathcal{M}$, an ambiguity attitudes parameter $\zeta \in [0, \infty]$, and a reference prior $\mu \in \Delta(\mathcal{M})$ such that $\succsim$ is represented by:*

$$\min_{\nu \in \Delta(\mathcal{M})} \left\{ \int_{\mathcal{M}} \min_{p \in \Delta(\Omega)} \{\mathbb{E}_p[u(f)] + \lambda_m R(p||m)\}\, d\nu(m) + \zeta R(\nu||\mu) \right\} \tag{11}$$

and $\min_p \{\mathbb{E}_p[u(f)] + \lambda_m R(p||m)\}$ represents $\succsim^m$ for all $m \in \mathcal{M}$.

[13] This shape of the cost κ can be obtained by requiring that preferences satisfy a stronger form of independence, namely the certainty independence axiom introduced by Gilboa and Schmeidler (1989). We omit the formal result for brevity.

Moreover, u is unique up to positive affine transformations, μ is unique, and $(\lambda_m)_{m\in\mathcal{M}}$ and ζ are unique given u.

In the representation (11), the relative entropy $R(p||m)$ measures the distance of p from model m and λ_m captures the degree of concern that model m may not be a good approximation of the true DGP. The lower λ_m is, the greater the DM's concern for misspecification. Similarly, the relative entropy $R(\nu||\mu)$ measures the distance of ν from the reference prior μ, and ζ captures the DM's aversion to ambiguity.

We now provide an axiomatic foundation for a Bayesian version of misspecification averse preferences whose functional form is also discussed in Cerreia-Vioglio et al. (2025). We introduce the following axiom.

Axiom 9 ($\mathcal{A}$-SEU):

(i) $\mathcal{A}$-Tradeoff Consistency. *For all nonnull events $E, A \in \mathcal{A}$, for all $x, y, z, w \in X$, and $f, g \in \mathcal{F}(\mathcal{A})$, if $xEf \succsim yEg$, $zEf \succsim wEg$, and $xAf \succsim yAg$, then $zAf \succsim wAg$.*

(ii) $\mathcal{A}$-S-Continuity. *For all finite partitions $\{E_1, \dots, E_n\} \subseteq \mathcal{A}$, for all $x, y \in X$, and $f \in \mathcal{F}_0(\mathcal{A})$, the sets $\{(\alpha_1, \dots, \alpha_n) \in [0,1]^n : \sum_{i=1}^n \chi_{E_i}[\alpha_i x + (1-\alpha_i)y] \succsim f\}$ and $\{(\alpha_1, \dots, \alpha_n) \in [0,1]^n : f \succsim \sum_{i=1}^n \chi_{E_i}[\alpha_i x + (1-\alpha_i)y]\}$ are closed.*

This axiom adapts to our setting the Tradeoff Consistency and S-Continuity axioms introduced by Wakker (2013) to axiomatize subjective expected utility on arbitrary state spaces. We then obtain the following characterization of the Bayesian aggregator.

Theorem 4: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map $\mathfrak{q}$ and $|\mathcal{M}| > 2$. The following statements are equivalent:*

(i) $\succsim$ is a misspecification averse preference relation satisfying Axioms 5, 6, and 9,

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

$$\int_{\mathcal{M}} \phi\left(\min_{p\in\Delta}\{\mathbb{E}_p[u(f)] + c(p,m)\}\right) d\mu(m) \tag{12}$$

and $\min_p\{\mathbb{E}_p[u(f)] + c(p,m)\}$ represents $\succsim^m$ for all $m \in \mathcal{M}$.

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 DM's concern about misspecification is reflected in the fact that, even after conditioning on the information that m is the best-fit model, she still takes into account other distributions that are sufficiently close to m. In this Bayesian specification, uncertainty about the identity of the best-fit model and the DM's attitudes toward this uncertainty, are captured, respectively, by the Bayesian prior μ over the set of hypothesized models and the index of ambiguity attitudes ϕ. The subjective belief μ quantifies which models the DM considers more likely to be good approximations of the true DGP. The curvature of ϕ captures the DM's attitudes toward ambiguity about the best-fit model: ϕ is concave (convex) if and only if the DM is ambiguity averse (ambiguity loving). The Bayesian criterion (12) can be viewed as an extension of the smooth ambiguity model of Klibanoff et al. (2005) that incorporates misspecification concerns. We recover the smooth ambiguity model by letting the misspecification aversion index c go to infinity off the diagonal (while remaining 0 on the diagonal), which corresponds to the limit case in which the DM is neutral with respect to misspecification. As already remarked in the introduction, this criterion specializes to the average robust control criterion axiomatized by Lanzani (2025) when the DM is neutral toward ambiguity about the identity of the best-fit model; this is precisely the case in which the index ϕ is affine. The relative entropy formulation of the misspecification aversion index $c(\cdot, m) = \lambda R(\cdot||m)$ can be obtained by imposing Savage's sure-thing principle, as discussed above, and by requiring the DM to display a uniform concern for misspecification across models in $\mathcal{M}$.

We conclude this section by showing that the cautious criterion axiomatized in Cerreia-Vioglio et al. (2025) arises as a special case of the representation in Theorem 2 when preferences exhibit a cautious attitude toward uncertainty about the identity of the best-fit model

Axiom 10 ($\mathcal{M}$-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 conceptual analogue, in our framework, of the caution axiom introduced by Gilboa et al. (2010) and used in Cerreia-Vioglio et al. (2025). The set of hypothesized models induces a (typically incomplete) dominance relation $\succsim_{\mathcal{M}}$ on $\mathcal{F}$, defined for all $f, g \in \mathcal{F}$ by:

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

If $f \succsim_{\mathcal{M}} g$, then f is better than g according to each model $m \in \mathcal{M}$, after taking into account misspecification concerns. Because the DM trusts the set of hypothesized models, whenever $f \succsim_{\mathcal{M}} g$ she is confident that f is better than g. Axiom 10 can then be rewritten as the requirement that if $f \not\succsim_{\mathcal{M}} x$, then $x \succsim f$. The interpretation is that if the DM is not sure that the uncertain act f is better than the constant act x, which is unaffected by uncertainty considerations, then she behaves cautiously and prefers the certain act to the uncertain one.

The following result shows that $\mathcal{M}$-Caution delivers, as a special case, the cautious criterion of Cerreia-Vioglio et al. (2025).

THEOREM 5: *Suppose that the triple* $(\Omega, \mathcal{G}, \mathcal{M})$ *admits a best-fit map* $\mathfrak{q}$. *The following statements are equivalent:*

(i) $\succsim$ *is a misspecification averse preference relation satisfying Axioms 5, 6, and 10,*

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

$$V(f) = \inf_{m \in \mathcal{M}} \min_{p \in \Delta} \{\mathbb{E}_p[u(f)] + c(p, m)\} \tag{13}$$

and $\min_p \{\mathbb{E}_p[u(f)] + c(p, m)\}$ *represents* $\succsim^m$ *for all* $m \in \mathcal{M}$.

In addition, if $\mathcal{M}$ *is compact, we have:*

$$V(f) = \min_{p \in \Delta} \{\mathbb{E}_p[u(f)] + \min_{m \in \mathcal{M}} c(p, m)\} \ . \tag{14}$$

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 models $m' \in \mathcal{M}$. Thus, the function $C_{\mathcal{M}}(\cdot) := \min_{m \in \mathcal{M}} c(\cdot, m)$ can be viewed as a statistical distance of probability distributions from the set of hypothesized models $\mathcal{M}$. It captures the DM's degree of misspecification aversion when she adopts a worst-case scenario approach to ambiguity about which model has the best fit.

6. Misspecification Aversion and Robust Monopoly Pricing

To illustrate how misspecification aversion can have qualitatively different implications from ambiguity aversion, we revisit the monopoly pricing example from Ball and Kattwinkel (2024). A seller (the DM) wants to sell a single good to a buyer. The seller does not know the buyer's valuation $\omega \in \Omega = [0, \bar{\theta}]$ for the good, but is informed that the buyer's median valuation is $\gamma \in (0, \bar{\theta})$. Thus, the set of probabilistic models $\mathcal{M}$ entertained by the seller consists of all probability distributions on Ω with median equal to γ. If the seller chooses a price $P \in \mathbb{R}_+$, only buyers whose valuation is at least P purchase the good. Hence, each price P induces an act $f_P(\omega) = P\mathbf{1}\{\omega \geq P\}$. The seller's payoff guarantee of a price $P \geq 0$ is given by the worst-case expected profit:

$$\inf_{m \in \mathcal{M}} \mathbb{E}_m[f_P] = \inf_{m \in \mathcal{M}} P \ m(\{\omega \geq P\}) \ .$$

The payoff guarantee of a price P is said to be robust if

$$\liminf_n \mathbb{E}_{q_n}[f_P] \geq \inf_{m \in \mathcal{M}} \mathbb{E}_m[f_P]$$

for all $m \in \mathcal{M}$ and sequences $(q_n)_n$ in $\Delta(\Omega)$ converging to m. Ball and Kattwinkel (2024) show that the maxmin optimal price, that is the price that maximizes the payoff guarantee, is $P^* = \gamma$, but that the payoff guarantee of P^* is not robust. We argue that a seller who is concerned about the possibility that the set $\mathcal{M}$ is misspecified would instead choose a price that delivers a robust payoff guarantee. Suppose that the seller's preferences are represented by the cautious misspecification averse criterion:

$$V^\lambda(f_P) = \inf_{m \in \mathcal{M}} \inf_{q \in \Delta(\Omega)} \{\mathbb{E}_q[f_P] + \lambda c_\eta(q, m)\}$$

with

$$c_\eta(q, m) = W_1(q, m) + \eta R(q||m)$$

where W_1 is the Wasserstein distance and $\eta > 0$.[14] As discussed in the previous section, we interpret $\lambda \in (0, \infty]$ as an index of misspecification aversion. The limit case $\lambda = \infty$ corresponds to the absence of concern for misspecification and thus to to the maximization of the payoff guarantee. We then obtain the following result.

Proposition 3: *Fix $\gamma \in (0, \bar{\theta})$ and $\eta < \gamma/2$. For all $\lambda \in (0, \infty)$, we have that the solution to $\max_{P \geq 0} V^\lambda(f_P)$ is given by*

$$\hat{P}(\lambda) = \frac{\lambda\gamma + \lambda\eta \log(1 + 2/\lambda)}{1 + \lambda} \in (0, \gamma) \ .$$

Moreover, $\partial\hat{P}(\lambda)/\partial\lambda > 0$ and

$$\lim_{\lambda \to 0} \hat{P}(\lambda) = 0 \quad \text{and} \quad \lim_{\lambda \to \infty} \hat{P}(\lambda) = \gamma.$$

In this case, when choosing the optimal price, a seller who is concerned about misspecification anticipates that the set of hypothesized models $\mathcal{M}$ may not contain the true distribution of the buyer's valuations and therefore chooses a price lower than the maxmin optimal one. It is immediate that this implies that the payoff guarantee of the price $\hat{P}(\lambda)$ is robust for all $\lambda < \infty$. Thus, even a minimal degree of concern about misspecification is enough to restore robustness of the payoff guarantee, while

[14] To obtain this result, we could simply choose $c(q, m) = W_1(q, m)$. The additional relative entropy perturbation is introduced only to ensure that c_η satisfies the properties required in the representation theorem.

still allowing the DM to obtain a guarantee arbitrarily close to that associated with the maxmin-optimal price under correct specification of the set of valuation distributions.

7. Conclusion

This paper provides an axiomatic foundation for a general class of misspecification averse preferences. We study a framework in which the DM formulates a possibly misspecified set of probabilistic models that she regards as plausible descriptions of the environment. We introduce the notion of a best-fit map that identifies, for each realization of the state of the world, the model in this set that best approximates the true DGP. This construction allows us to disentangle the DM's concern that the set of models may be misspecified from her attitudes toward ambiguity about which hypothesized model is the most accurate description of the environment. Our main result shows that the DM's preferences can be represented as a monotone aggregation of misspecification averse evaluations based on each model. As shown in the paper, this representation separates attitudes toward ambiguity, captured by the aggregator, from misspecification concerns, captured by the misspecification averse model-conditional evaluations. Specific functional forms of the aggregator can be obtained by imposing additional behavioral axioms on the DM's preferences. In particular, we show that the important decision criteria recently introduced by Cerreia-Vioglio et al. (2025) and Lanzani (2025) fall within the general class of misspecification averse preferences analyzed here, and we provide axioms that deliver the Bayesian aggregator and the cautious criterion as special cases.

In this paper, we have focused on a DM facing a static decision problem. In many economic scenarios, however, agents face dynamic problems that involve making multiple decisions over time and reacting to new information. We address this issue in Maselli (2025), where we study a DM who employs a set of probabilistic models to solve an infinite-horizon problem under uncertainty and learns over time. There, we extend the analysis of misspecification averse preferences to a dynamic context and provide axiomatic foundations for forward-looking versions of the decision criteria discussed in this paper. In the context of the forward-looking analogue of the Bayesian aggregator, a satisfactory theory of dynamic decision making should also offer guidance on how the DM revises her beliefs in response to the arrival of new information. The traditional approach in the literature has been to require dynamic consistency. However, it is well documented that under uncertainty aversion, dynamic consistency is incompatible with Bayesian updating. At the same time, Bayes' rule has long established itself as the central procedure for belief revision, both because of its intuitive appeal and its desirable statistical properties. We contribute to this discussion by proposing a novel behavioral axiom that justifies updating the DM's beliefs via Bayes' rule even in the presence of ambiguity and misspecification concerns.

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.

Arrow, K. J. (1970): *Essays in the theory of risk-bearing*, North-Holland Amsterdam.

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

Ball, I. and D. Kattwinkel (2024): "Robust Robustness," *arXiv preprint arXiv:2408.16898*.

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.

Bonaglia, E. and P. Dedola (2025): "Foundations of Misspecification Aversion," *Mimeo*.

Box, G. E. (1976): "Science and statistics," *Journal of the American Statistical Association*, 71, 791–799.

——— (1979): "Robustness in the strategy of scientific model building," in *Robustness in statistics*, Elsevier, 201–236.

Cerreia-Vioglio, S., L. P. Hansen, F. Maccheroni, and M. Marinacci (2025): "Making decisions under model misspecification," *Review of Economic Studies*, rdaf046.

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.

Cox, D. R. (1995): "Discussion of the Paper by Chatfield," *Journal of the Royal Statistical Society Series A: Statistics in Society*, 158, 455–456.

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. 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.

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

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.

Kechris, A. (2012): *Classical descriptive set theory*, vol. 156, Springer Science & Business Media.

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. (2025): "Dynamic concern for misspecification," *Econometrica*, 93, 1333–1370.

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

Maselli, A. (2025): "Forward-Looking Preferences and Misspecification Aversion," *Mimeo*.

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 (2024): "Monotone additive statistics," *Econometrica*, 92, 995–1031.

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.

Wakker, P. P. (2013): *Additive representations of preferences: A new foundation of decision analysis*, vol. 4, Springer Science & Business Media.

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.

White, H. (1982): "Maximum likelihood estimation of misspecified models," *Econometrica: Journal of the econometric society*, 1–25.

Appendix

The appendices are organized as follows. Appendix A collects precise definitions of all the mathematical notions used in the paper and may be skipped on a first reading. Appendix B includes auxiliary mathematical results and proves that the model-conditional preferences are well-defined and inherit the basic axioms imposed on the overall preference relation $\succsim$. Appendix C is central: it shows that, under consistency, the representation of preferences is a monotone aggregation of the model-conditional certainty equivalents. It also shows that axioms imposed on the restriction of preferences to $\mathcal{F}(\mathcal{A})$ correspond to functional properties of the aggregator. Appendix D contains the proofs of the general representation results (Theorems 1 and 2) and of the comparative statics characterizations. Appendix E collects the proofs of the special cases of the aggregator. Appendix F contains the proofs of the results for the application to the robust monopoly pricing example.

Appendix A. Mathematical Preliminaries

A.1. **Basic Notions.** Given an arbitrary measurable space $(Y, \mathcal{Y})$, we denote by $\Delta(Y, \mathcal{Y})$ the space of countably additive probability measures on $(Y, \mathcal{Y})$. Sometimes, we will omit making explicit reference to the sigma-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 $\Delta(Y, \mathcal{Y})$ with the Borel sigma-algebra generated by this topology; which is the same as the natural sigma-algebra $\mathcal{D}^{Y,\mathcal{Y}}$ generated by the family of evaluation maps:

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

and any subset $\mathcal{Q}$ of Δ, with the relative sigma-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 $\Psi(\varphi) \geq \Psi(\varphi')$.[15] 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 $\Psi(k\chi_{E_{n_0}} + \varphi\chi_{E^c_{n_0}}) > \Psi(\varphi')$.

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 sigma-algebra generated by f.

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, $\operatorname{dom}\varphi := \{p \in C : |\varphi(p)| < \infty\}$. Moreover, we say such function φ to be *grounded* if $\inf_{p \in C} \varphi(p) = 0$. Fix a subset $\mathcal{Q} \subseteq \Delta(Y, \mathcal{Y})$ of countably additive probability measures. We say that a function $c : \Delta(Y, \mathcal{Y}) \times \mathcal{Q} \to [0, \infty]$ is a *statistical divergence distance* if it satisfies the following properties:

(i) for each $q \in \mathcal{M}$, $p = q$ if and only if $c(p, q) = 0$,
(ii) c is jointly lower semicontinuous,
(iii) for all $q \in \mathcal{Q}$, $p \in \operatorname{dom} c(\cdot, q)$ implies that $p \ll q$.

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

Furthermore, a statistical divergence distance c is convex if the section $c(\cdot, q)$ is a convex function for each $q \in \mathcal{Q}$

A.3. **Statistical Framework.** A maintained assumption in all the paper is that there is triple $(\Omega, \mathcal{G}, \mathcal{M})$, where $(\Omega, \mathcal{G})$ is a measure space with countably generated sigma-algebra $\mathcal{G}$ and $\mathcal{M} \subseteq \Delta(\mathcal{G}) := \Delta(\Omega, \mathcal{G})$ is a set of models. When this triple admits a best-fit map $\mathfrak{q}$, we always denote by $\mathcal{A}$ the sigma-algebra generated by $\mathfrak{q}$ and set $E^m := \mathfrak{q}^{-1}(m)$ all $m \in \mathcal{M}$. Moreover, we denote by $\mathcal{D} := \mathcal{D}^{\Omega,\mathcal{G}}$ and $\mathcal{D}_{\mathcal{M}} := \mathcal{D}^{\Omega,\mathcal{G}}_{\mathcal{M}}$ respectively the natural sigma-algebra on $\Delta(\mathcal{G})$ and the relative sigma-algebra on $\mathcal{M}$.

Appendix B. Auxiliary Results

Denote by Λ the set of all the events in $\mathcal{G}$ that have probability either 0 or 1 according to all models $m \in \mathcal{M}$:

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

Lemma B.1: *Suppose that the triple $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map. The sigma-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 model $m \in \mathcal{M}$.*

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

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

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

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

and, therefore, $\mathfrak{q}^{-1}\left(\{p \in \Delta(\mathcal{G}) : p(E) \leq r\}\right) \in \Lambda$, 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 sigma-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-sigma-algebra of $\mathcal{G}$, we can define the operator $u : \mathcal{F}(\mathcal{E}) \to B(\mathcal{E})$ 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, $u : \mathcal{F}(\mathcal{E}) \to B(\mathcal{E})$ is an affine operator. Moreover, $\{u(f) : f \in \mathcal{F}_0(\mathcal{E})\} = B_0(\mathcal{E})$ and $\{u(f) : f \in \mathcal{F}(\mathcal{E})\} = B(\mathcal{E})$.*

PROOF: Take any $f \in \mathcal{F}_0(\mathcal{E})$. 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 $\varphi \in \mathcal{F}_0(\mathcal{E})$. This shows that $B_0(\mathcal{E}) \subseteq \{u(f) : f \in \mathcal{F}_0(\mathcal{E})\}$.

Take now $f \in \mathcal{F}$. Then, we can find a finite set $X_0 \subseteq X$ such that $\operatorname{Im} f \subseteq \operatorname{co} X_0$. Moreover, the latter set is compact and u is continuous on co_X since it is affine. It follows that $u(f)$ is measurable and $\min \operatorname{co} X_0 \leq u(f(\omega)) \leq \max \operatorname{co} X_0$ for all $\omega \in \Omega$. We conclude that $u(f) \in B(\mathcal{E})$. As for the other direction, take $\varphi \in B(\mathcal{E})$. Then, we can find $k < K$ such that $k \leq \varphi \leq K$. Since u is surjective, we can pick $x_k, x_K \in X$ such that $u(x_k) = k$ and $u(x_K) = K$. Then, define the function $\alpha_\varphi : \Omega \to [0,1]$ as

$$\alpha_\varphi(\omega) = \frac{\varphi(\omega) - k}{K - k}$$

and notice that it is also measurable. Define $f_\varphi(\omega) = x_k + \alpha_\varphi(\omega)(x_K - x_k)$ and notice that it is also measurable. Finally, using affinity of u, we obtain that for all $\omega \in \Omega$,

$$u(f_\varphi) = u(x_k) + \alpha(\omega)_\varphi(u(x_K) - u(x_k)) = k + \frac{\varphi(\omega) - k}{K - k}(K - k) = \varphi(\omega).$$

We conclude that $B(\mathcal{E}) \subseteq \{u(f) : f \in \mathcal{F}(\mathcal{E})\}$. ■

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 $f \in \mathcal{F}$. By definition, we can find a finite set $X_0 \subseteq X$ such that $\operatorname{Im} f \subseteq \operatorname{co} X_0$. Then, we can pick x^* and x_* in X_0 such that $x^* \succsim x \succsim x_*$ for all $x \in X_0$. Axiom 1 then implies 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 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 two axioms on a weak order on $\mathcal{F}$.

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

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

AXIOM B.2 (Weak Certainty Independence): *For all $f, f' \in \mathcal{F}$, $x, y \in X$, and $\alpha \in (0,1)$,*

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

LEMMA B.4: *Suppose that $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map and that the preference relation $\succsim$ satisfies Axioms 1, 2, 3, and 4. For all $m \in \mathcal{M}$, define $\succsim^m$ as follows: for all $f, f' \in \mathcal{F}$,*

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

Then, $\succsim^m$ is well-defined, satisfies Axioms 1 and B.1, and coincides with $\succsim$ when restricted to constant acts in X. Moreover, $\succsim^m$ satisfies Axiom B.2 if $\succsim$ satisfies Axiom 6, and $\succsim^m$ satisfies Axiom 5 if $\succsim$ satisfies Axiom 5.

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 2 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. $\succsim^m|_X = \succsim_X$. By what is shown above, for all $x, y \in X$, $x \succsim^m y$ if and only if $x E^m h \succsim y E^m h$ for all $h \in \mathcal{F}$. Take $g = x$ and appy Axiom 2 to obtain that

$$x \succsim^m y \iff x \succsim y E^m x \iff x \succsim y$$

proving the result.

Step 4: Non-triviality. This follows immediately by Step 3.

Step 5: 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 $f E^m g \succsim f' E^m g$ and, therefore, $f \succsim^m f'$, proving the statement.

Step 6: Unboundedness. This follows immediately by Step 3.

Step 7. 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 3, we have

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

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

Step 8: 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 $E_n \coloneqq A_n \cap E^m$ 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 5, 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.

Step 9: Weak Certainty Independence. Suppose that $\succsim$ satisfies Axiom 6. 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 6,

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

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. ■

Appendix C. From Acts to Second-Order Acts

Throughout this section, assume that the triple $(\Omega, \mathcal{G})$ admits a best-fit map $\mathfrak{q}$ generating the sigma-algebra $\mathcal{A}$, and that there exists a family of operators $I^m : B(\mathcal{G}) \to \mathbb{R}$ for each $m \in \mathcal{M}$, satisfying the following properties:

(i) I^m is normalized, monotone, and continuous for all $m \in \mathcal{M}$;
(ii) for all $m \in \mathcal{M}$, $\varphi = \psi$ a.e. $[m]$ implies $I^m(\varphi) = I^m(\psi)$ for all $\varphi, \psi \in B(\mathcal{G})$;
(iii) for all $\varphi \in B(\mathcal{G})$, $m \mapsto I^m(\varphi)$ is measurable.

We begin by showing that for all $m \in \mathcal{M}$, I^m is linear when restricted to $B(\mathcal{A})$.

Lemma B.5: *For all $m \in \mathcal{M}$, the restriction of I^m to $B(\mathcal{A})$ is linear. In particular, I^m is Lipschitz continuous of order 1 on $B(\mathcal{A})$.*

Proof of Lemma B.5: Fix $m \in \mathcal{M}$ arbitrarily. First of all, we show that I^m is additive on $B(\mathcal{A})$. Since $B_0(\mathcal{A})$ is dense in $B(\mathcal{A})$, we can find sequences $(\varphi_n)_{n\in\mathbb{N}}$, $(\psi_n)_{n\in\mathbb{N}} \subseteq B_0(\mathcal{A})$ such that $\varphi_n \to \varphi$ and $\psi_n \to \psi$ in the supnorm. Fix any $n \in \mathbb{N}$. Then, there exists a partition $(E_i^n)_{i=1}^{k_n} \subseteq \mathcal{A}$ such that $\varphi_n = \sum_{i=1}^{k_n} r_i^n \chi_{E_i^n}$ and $\psi_n = \sum_{i=1}^{k_n} \tilde{r}_i^n \chi_{E_i^n}$ for reals $(r_i^n)_{i=1}^{k_n}, (\tilde{r}_i^n)_{i=1}^{k_n} \subseteq \mathbb{R}$. By Lemma B.1, for each $m \in \mathcal{M}$, there exists a unique $j_n(m) \in \{1, \ldots, k_n\}$ such that $m(E_{j_n(m)}^n) = 1$ and $m(E_i^n) = 0$ for all $i \neq j_n(m)$. Therefore, for all $m \in \mathcal{M}$, $\varphi_n = r_{j_n(m)}^n$ and $\psi_n = \tilde{r}_{j(m)}^n$ a.e.$[m]$ and, similarly $\varphi_n + \psi_n = r_{j_n(m)}^n + \tilde{r}_{j_n(m)}^n$, so that

$$\begin{aligned} I^m(\varphi_n + \psi_n) &= I^m(r_{j_n(m)}^n + \tilde{r}_{j_n(m)}^n) \\ &= r_{j_n(m)}^n + \tilde{r}_{j_n(m)}^n \\ &= I^m(r_{j_n(m)}^n) + I^m(\tilde{r}_{j_n(m)}^n) = I^m(\varphi_n) + I^m(\psi^n). \end{aligned}$$

We conclude that $I^m(\varphi_n + \psi_n) = I^m(\varphi_n) + I^m(\psi_n)$ for all $n \in \mathbb{N}$. Since I^m is continuous with respect to supnorm convergence, taking limits, we conclude that $I^m(\varphi + \psi) = I^m(\varphi) + I^m(\psi)$.

We now show that I^m is homogeneous. Take $\varphi \in B(\mathcal{A})$ and $\kappa \in \mathbb{R}$. As before, we can find a sequence $(\varphi_n)_{n\in\mathbb{N}} \subseteq B_0(\mathcal{A})$ such that $\varphi_n \to \varphi$ in the supnorm. Notice also that $||\kappa\varphi_n - \kappa\varphi||_\infty = |\kappa|||\varphi_n - \varphi||_\infty \to 0$. Fix any n and pick a partition $(E_i^n)_{i=1}^{k_n} \subseteq A$ such that $\varphi_n = \sum_{i=1}^{k_n} r_i^n \chi_{E_i^n}$ for reals $(r_i^n)_{i=1}^{k_n} \subseteq \mathbb{R}$. For each $m \in \mathcal{M}$, Lemma B.1 implies that there exists a unique $j_n(m) \in \{1, \ldots, k_n\}$ such that $m(E_{j_n(m)}^n) = 1$ and $m(E_i^n) = 0$ for all $i \neq j_n(m)$. Therefore, for all $m \in \mathcal{M}$, $\varphi_n = r_{j_n(m)}^n$ a.e. $[m]$ and, similarly $\kappa\varphi_n = \kappa r_{j_n(m)}^n$ a.e. $[m]$, so that

$$I^m(\kappa\varphi_n) = I^m(\kappa r_{j_n(m)}^n) = \kappa r_{j_n(m)}^n = \kappa I^m(r_{j_n(m)}^n) = \kappa I^m(\varphi_n)\ .$$

Therefore, $I^m(\kappa\varphi_n) = \kappa I^m(\varphi_n)$ for all $n \in \mathbb{N}$ and taking limits and by continuity of I^m we conclude that $I^m(\kappa\varphi) = \kappa I^m(\varphi)$.

■

The following lemma shows that we are able to find a non-linear conditional expectation given $\mathcal{A}$ that is common to all hypothesized models $m \in \mathcal{M}$.

Lemma B.6: *The map $m \mapsto I^m(\varphi)$ is bounded and there exists a non-linear 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 $\omega \in \Omega$ and for all*

$A \in \mathcal{A}$ and $m \in \mathcal{M}$,

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

Proof of Lemma B.6: First, we show that $m \mapsto I^m(\varphi)$ is bounded. Indeed, take $\varphi \in B(\Omega, \mathcal{G})$. Then, there exist $k, K \in \mathbb{R}$ such that $k \leq \varphi \leq K$. Since for each $m \in \mathcal{M}$, I^m is normalized and monotone, we have that

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

proving boundedness.

Fix any $\varphi \in B(\Omega, \mathcal{G})$. Since $m \mapsto I^m(\varphi)$ is bounded and $\mathcal{D}_{\mathcal{M}}$-measurable, it follows that the composition

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

is a $\mathcal{A}$-measurable and bounded functional. Obtain $I_{\mathcal{A}}(\varphi)$ by defining $I_{\mathcal{A}}(\varphi)(\omega) = I^{\mathfrak{q}(\omega)}\varphi$ for all $\omega \in \Omega$. 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 $E^m = \{\omega \in \Omega : \mathfrak{q}(\omega) = m\}$. 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

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

Now, notice that $I_{\mathcal{A}}(\varphi)(\omega) = I^m(\varphi)$ for all $\omega \in E^m$ and, therefore, $I_{\mathcal{A}}(\varphi)\chi_A = I^m(\varphi)$ a.e. $[m]$ if $m(A) = 1$ and $I_{\mathcal{A}}(\varphi)\chi_A = 0$ a.e. $[m]$ if $m(A) = 0$. Similarly, $\varphi\chi_A = \varphi$ a.e. $[m]$ if $m(A) = 1$ and $\varphi\chi_A = 0$ a.e. $[m]$ if $m(A) = 0$. Then, by property *(ii)* of I^m, it follows that

$$\begin{aligned} m(A) = 1 &\implies I^m(I_{\mathcal{A}}(\varphi)\chi_A) = I^m(I^m(\varphi)) = I^m(\varphi) = I^m(\varphi\chi_A) \\ m(A) = 0 &\implies I^m(I_{\mathcal{A}}(\varphi)\chi_A) = I^m(0) = 0 = I^m(\varphi\chi_A)\ . \end{aligned}$$

.

■

Notice that for each $\varphi \in B(\mathcal{G})$, we can see $I^m(\varphi)$ as a function from 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.6, we have that $\operatorname{Im} T \subseteq B(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$. Moreover, we have the following result.

LEMMA B.7: *$T : B(\Omega, \mathcal{G}) \to B(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$ is Lipschitz continuous of order 1 with respect to supnorm convergence and is additive and homogeneous on $B(\mathcal{A})$.*

PROOF OF LEMMA B.7: Take $\varphi, \psi \in B(\mathcal{A})$ and $k \in \mathbb{R}$. By Lemma B.5, I^m is linear when restricted to $B(\mathcal{A})$ for all $m \in \mathcal{M}$. Therefore:

$$T(\varphi + \kappa\psi)(m) = I^m(\varphi + \kappa\psi) = I^m(\varphi) + \kappa I^m(\psi) = T(\varphi)(m) + \kappa T(\psi)(m)$$

for all $m \in \mathcal{M}$. Thus, $T(\varphi + \kappa\psi) = T(\varphi) + \kappa T(\psi)$, showing linearity. Moreover, since each I^m is Lipschitz continuous of order 1 on $B(\mathcal{A})$, we have that:

$$|T(\varphi)(m) - T(\psi)(m)| = |I^m(\varphi) - I^m(\psi)| \leq ||\varphi - \psi||_\infty$$

and, therefore,

$$||T(\varphi) - T(\psi)||_\infty = \sup_{m \in \mathcal{M}} |T(\varphi)(m) - T(\psi)(m)| \leq ||\varphi - \psi||_\infty$$

showing that T is also Lipschitz continuous of order 1 on $B(\mathcal{A})$.

■

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-sigma-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 Lemma B.6, $I_{\mathcal{A}}(\varphi_\xi) \in B(\mathcal{A})$ 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, by Lemma B.7, T is Lipschitz of order 1 and, therefore

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

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 $I^m(\varphi_n) = I^m(r_n^l) = r_n^l$ 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 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* $T(\{\chi_E : E \in \mathcal{A}\}) = \{\chi_D : D \in \mathcal{D}_\mathcal{M}\}$. *Moreover, for each* $\xi \in B_0(\mathcal{M}, \mathcal{D}_\mathcal{M})$ *such that* $0 \leq \xi \leq 1$, *there is* $\varphi \in B(\mathcal{A})$ *with* $0 \leq \varphi \leq 1$ *such that* $\xi = T(\varphi)$.

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, therefore, 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, $D_E \coloneqq [T(\chi_E)]^{-1}(\{1\}) \in \mathcal{D}_{\mathcal{M}}$ 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 $E^D = \mathfrak{q}^{-1}(D)$. Then, $E^D \in \mathcal{A}$ 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 $\xi = T(\varphi_\xi) = \sum_{i=1}^k \chi_{D_{E_i}} r_i \in B_0(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$. □

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 $i = 1, \ldots, k$, we can find $E^{D_i} \in \mathcal{A}$ such that $\chi_{D_i} = T(\chi_{E^{D_i}})$. Define $\varphi_\xi \coloneqq \sum_{i=1}^k \chi_{E^{D_i}} r_i$. 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, $T(\chi_{E^{D_{j_m}}})(m) = I^m(\chi_{E^{D_{j_m}}}) = I^m(0) = 0 \neq \chi_{D_{j_m}}(m) = 1$, 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 $\operatorname{Im} T = B(\mathcal{M}, \mathcal{D}_\mathcal{M})$.

The last part of the result follows by steps *iii* and *iv* and by Lemma B.7. ■

PROPOSITION B.1: *The following are equivalent:*

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

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

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

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

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.*
- $\hat{I}$ *is quasiconcave if and only if* I *is quasiconcave.*
- $\hat{I}$ *is translation invariant if and only if* I *is translation invariant.*

PROOF OF PROPOSITION B.1:
(i) implies (ii). Define $\hat{I} : B(\mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ as follows: for all $\xi \in B(\mathcal{D}_{\mathcal{M}})$,

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

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

Step 1: $\hat{I}$ is well-defined. Pick $\xi \in B(\mathcal{D}_{\mathcal{M}})$ arbitrarily. That a $\varphi_\xi \in B(\mathcal{A})$ such that $\xi = T(\varphi_\xi)$ exists follows from Lemma B.9. Moreover, suppose there are two $\varphi, \psi \in B(\mathcal{A})$ 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 $\xi, \xi' \in B(\mathcal{D}_{\mathcal{M}})$ such that $\xi \geq \xi'$. By Lemma B.9, we can find $\varphi_\xi, \varphi_{\xi'} \in B(\mathcal{A})$ 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 $\tilde{I} : B(\mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ such that $I(\varphi) = \tilde{I}(T(\varphi))$ for all $\varphi \in B(\mathcal{A})$. Then, take any $\xi \in B(\mathcal{D}_{\mathcal{M}})$. By Lemma B.9, there exists $\varphi_\xi \in B(\mathcal{A})$ 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 $\xi, \xi' \in B(\mathcal{D}_{\mathcal{M}})$ and $c \in \mathbb{R}$. First we show that the set $L = \{\alpha \in [0,1] : \hat{I}(\alpha\xi + (1-\alpha)\xi') \leq c\}$ is closed. If it is empty, it is closed. If it is nonempty, take any sequence $(\alpha_n)_n$ in L such that $\alpha_n \to \alpha_0$. By Lemma B.9, we can pick $\varphi, \varphi' \in B(\mathcal{A})$ such that $\xi = T(\varphi)$ and $\xi' = T(\varphi')$. Lemma B.5 implies that

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

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. Since this holds for all $\xi, \xi' \in B(\mathcal{D}_\mathcal{M})$ 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 $U_c = \{\xi \in B(\mathcal{D}_\mathcal{M}) : \hat{I}(\xi) \geq c\}$ is convex. If it is empty, this is 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 $\varphi_1, \varphi_2 \in B(\mathcal{A})$ 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 $I(\varphi_2) = \hat{I}(\xi_2) \geq c$. Since I is quasiconcave, it follows that $I(\alpha\varphi_1 + (1-\alpha)\varphi_2) \geq c$. Moreover, Lemma B.5 implies that $T(\alpha\varphi_1 + (1-\alpha)\varphi_2) = \alpha\xi_1 + (1-\alpha)\xi_2$. Then:

$$\hat{I}(\alpha\xi_1 + (1-\alpha)\xi_2) = I(\alpha\varphi_1 + (1-\alpha)\varphi_2) \geq c$$

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 $\xi, \xi' \in B(\mathcal{D}_\mathcal{M})$ 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 $\varphi, \varphi' \in B(\mathcal{A})$ such that $\xi = T(\varphi)$ and $T(\varphi') = \xi'$. It follows that $I(\varphi) = \hat{I}(\xi) > \hat{I}(\xi') = I(\varphi')$. Let $E_n \coloneqq \mathfrak{q}^{-1}(D_n) \in \mathcal{A}$ 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

$$\begin{aligned} m(E_{n_0}) = 1 &\implies kE_{n_0}\varphi = k \text{ } m\text{-a.e.} \implies I^m(kE_{n_0}\varphi) = I^m(k) = k \\ m(E_{n_0}) = 0 &\implies kE_{n_0}\varphi = \varphi \text{ } m\text{-a.e.} \implies I^m(kE_{n_0}\varphi) = I^m(\varphi) = \xi(m) \end{aligned}$$

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.

Step 8: $\hat{I}$ is translation invariant. Take $\xi \in B(\mathcal{D}_\mathcal{M})$ and $k \in \mathbb{R}$. Then, we can find $\varphi \in B(\mathcal{A})$ such that $\xi = T(\varphi)$. Since I is translation invariant and normalized, $I(\varphi + k) = I(\varphi) + k$. On the other hand, since each I^m is translation invariant,

$I^m(\varphi + k) = I^m(\varphi) + k$ for all $m \in \mathcal{M}$ and, thus, $T(\varphi + k) = T(\varphi) + k = \xi + k$. Then:

$$\hat{I}(\xi + k) = \hat{I}(T(\varphi + k)) = I(\varphi + k) = I(\varphi) + k = \hat{I}(\xi) + k$$

as we wanted to show.

(ii) implies (i).

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

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 $\varphi, \varphi' \in B(\mathcal{A})$ 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 $\varphi, \varphi' \in B(\mathcal{A})$ and $I^m(\varphi) \geq I^m(\varphi')$ for all $m \in \mathcal{M}$, then $I(\varphi) \geq I(\varphi')$.

Take any two $\varphi, \varphi' \in B(\mathcal{A})$ 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 $(\varphi_n)_n$ in $B(\mathcal{A})$ such that $\varphi_n \to \varphi \in B(\mathcal{A})$ 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, $T(\varphi_n), T(\varphi) \in B(\mathcal{D}_{\mathcal{M}})$. 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 $\varphi_1, \varphi_2 \in B(\mathcal{A})$ 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 continuous. Take $\varphi, \varphi' \in B(\mathcal{A})$ 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 $m \notin D_n$ 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 $m \in \mathcal{M}$,

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

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

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

as we wanted to show.

Step 7: I is translation invariant. Take $\varphi \in B(\mathcal{A})$ and $k \in \mathbb{R}$. Since I^m is translation invariant, $I^m(\varphi + k) = I^m(\varphi) + k$ for all $m \in \mathcal{M}$. That is, $T(\varphi + k) = T(\varphi) + k$. Since $\hat{I}$ is translation invarian and normalized,

$$I(\varphi + k) = \hat{I}(T(\varphi + k)) = \hat{I}(T(\varphi) + k) = \hat{I}(T(\varphi)) + k = I(\varphi) + k$$

as we wanted to show. ■

Appendix D. Proofs of the Results in Section 4

Proof of Theorem 1: *(i) implies (ii).*

Suppose that $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map and the preference relation $\succsim$ satisfies Axioms 1, 2, 3, and 4. Since $\succsim$ is a non-trivial, algebraically 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 $m \in \mathcal{M}$ and define $\succsim^m$ as in Lemma B.4. The lemma implies that $\succsim^m |_X = \succsim|_X$. Therefore, $\succsim^m$ is represented by u when restricted to constant acts in X. Define the functional $I^m : B(\mathcal{G}) \to \mathbb{R}$ as follows: for each $\varphi \in B(\mathcal{G})$, $I^m(\varphi) = u(x_{f_\varphi})$ 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 $f, f' \in \mathcal{F}$:

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

We first show that I^m is monotone, normalized, continuous, and quasiconcave.

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:

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

showing that I^m is normalized.

Monotonicity. Take $\varphi, \psi \in B(\mathcal{G})$ 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

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

which proves the claim.

Quasiconcavity. Take any $\varphi, \psi \in B(\mathcal{G})$ such that $I^m(\varphi) = I^m(\psi)$ 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 uncertainty aversion by Axiom 3, we have that

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

and, therefore:

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

proving the claim.

Continuity. It follows by a routine argument.

$\varphi = \psi$ *a.e.* $[m]$ *implies* $I^m(\varphi) = I^m(\psi)$.

Take $\varphi, \psi \in B(\mathcal{G})$ and assume that $\varphi = \psi$ a.e. $[m]$. Then, we can find a set E with $m(E) = 1$ such that $\varphi(\omega) = \psi(\omega)$ for all $\omega \in E$. Take $f_\varphi, f_\psi \in \mathcal{F}$ such that $u(f_\varphi) = \varphi$ and $u(f_\psi) = \psi$. Let $\tilde{f}_\psi = f_\varphi E f_\psi$. It is clear that $u(\tilde{f}_\psi) = \psi$ and $f_\varphi(\omega) = \tilde{f}_\psi(\omega)$ for all $\omega \in E$. Axiom 2, then, implies that $f_\varphi \sim^m \tilde{f}_\psi$ and we condlude that

$$I^m(\varphi) = I^m(u(f_\varphi)) = I^m(u(\tilde{f}_\psi)) = I(\psi)$$

as we wanted to show.

$m \mapsto I^m(\varphi)$ is lower semicontinuous for all $\varphi \in B(\mathcal{G})$. 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:

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

and the latter is closed since $\succsim$ satisfies Coherence.

We already know that $\succsim$ is represented by u when restricted to constant acts. Define the functional $I : B(\mathcal{G}) \to \mathbb{R}$ such that for each $\varphi \in B(\mathcal{G})$, $I(\varphi) := u(x_{f_\varphi})$, 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 $\varphi \in B(\mathcal{G})$, while the certainty equivalent $x_{f_\varphi} \sim f_\varphi$ exists by Lemma B.3. Moreover, for any $\varphi \in B(\mathcal{G})$, 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 $B(\mathcal{G})$. It is easily seen that such functional is also normalized, monotone, and continuous.[16]

Define the function $V := I \circ u : \mathcal{F} \to \mathbb{R}$. 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, let $I_{\mathcal{A}}$ be the generalized conditional expectation as in Lemma B.6. Take now $\varphi, \psi \in B(\mathcal{G})$ 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:

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

[16] See for example the proof of Theorem 1 (Omnibus) in the working paper version of Cerreia-Vioglio, Maccheroni, and Marinacci (2022).

By this fact and since I is monotone, normalized, and continuous, by Proposition B.1, there exists a unique monotone, normalized, continuous functional $\hat{I} : B(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ such that $I(\varphi) = \hat{I}(T(\varphi))$ for all $\varphi \in B(\mathcal{A})$. Take now any $\varphi \in B(\mathcal{G})$. Since $I_{\mathcal{A}}(\varphi) \in B(\mathcal{A})$, we also know that $I^m(I_{\mathcal{A}}(\varphi)) = I^m(\varphi)$ for all $m \in \mathcal{M}$. By what shown above, $I(\varphi) = I(I_{\mathcal{A}}(\varphi))$ and, therefore,

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

By letting $I(\cdot, m) = I^m$ for all $m \in \mathcal{M}$, we conclude that for all $f, g \in \mathcal{F}$,

$$f \succsim g \iff I(u(f)) \geq I(u(g)) \iff \hat{I}(I(u(f), \cdot)) \geq \hat{I}(I(u(g), m)).$$

Uniqueness follows by Proposition B.1 and routine arguments.

(ii) implies (i). It is clear that $\succsim$ satisfies Axiom 1. Fix any $m \in \mathcal{M}$. For any $f, g, h, h' \in \mathcal{F}$,

$$fE^m h \succsim gE^m h \iff I(u(f), m) \geq I(u(g), m) \iff fE^m h' \succsim gE^m h'.$$

Moreover, if $f = g$ a.e. $[m]$, then $u(f) = u(g)$ a.e. $[m]$ and, therefore, $I(u(f), m) = I(u(g), m)$, showing that $fE^m h \sim gE^m h$. Finally, for any $f \in \mathcal{F}$ and $x \in X$,

$$\{m \in \mathcal{M} : x \succsim fE^m x\} = \{m \in \mathcal{M} : I(u(f), m) \leq u(x)\} \tag{15}$$

and the latter set is closed since the map $m \mapsto I(u(f), m)$ is lower semicontinuous. It follows that $\succsim$ satisfies Axiom 2. Now, for each $m \in \mathcal{M}$, take $f, g \in \mathcal{F}$ such that $f \sim gE^m h$. It follows that $I(u(f), m) = I(u(g), m)$. Then, since $I(\cdot, m)$ is quasiconcave,

$$I(u(\alpha f + (1-\alpha)g), m) = I(\alpha u(f) + (1-\alpha)u(g), m) \geq I(u(f), m,) \tag{16}$$

and, therefore,

$$\alpha f + (1-\alpha)gE^m f = (\alpha f + (1-\alpha)g)E^m f \succsim fE^m f = f$$

showing that $\succsim$ satisfies Axiom 3. As for consistency, take acts $f, g, h \in \mathcal{F}$ and assume that $fE^m h \succsim gE^m h$ for all $m \in \mathcal{M}$. Then, $I(u(f), m) \geq I(u(g), m)$ for all $m \in \mathcal{M}$ and, since $\hat{I}$ is monotone, we conclude that $\hat{I}(u(f), \cdot) \geq \hat{I}(u(g), \cdot)$. Hence, $f \succsim g$, showing that $\succsim$ satisfies Axiom 4. We conclude that $\succsim$ is a misspecification averse preference.

■

Proof of Proposition 1: Suppose that $\succsim_1$ and $\succsim_2$ are two misspecification averse preferences represented respectively by $(\hat{I}_1, u_1, I_1)$ and $(\hat{I}_2, u_2, I_2)$ as in Theorem 1.

Point (i). Suppose that u_1 is a positive affine transformation of u_2 and that $I_1(\cdot, m) \leq I_2(\cdot, m)$ for all $m \in \mathcal{M}$. 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, $I_1(u(f), m) \geq u(x)$. Then:

$$I_2(u(f), m) \geq I_1(u(f), m) \geq u(x)$$

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 non-triviality 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 $\varphi \in B(\mathcal{G})$. 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 $\varphi \in B(\mathcal{G})$.

Point (ii). 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 normalization 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 (8) and non-triviality automatically imply that u_2 is a positive affine transformation of u_1. Assume that $u_1 = u_2 = u$ and take $\xi \in B(\mathcal{D}_{\mathcal{M}})$. 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 (8) 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, $\hat{I}_1(\xi) \leq \hat{I}_2(\xi)$ for all $\xi \in B(\mathcal{D}_{\mathcal{M}})$.

■

Proof of Theorem 2: *(i) implies (ii).*

Suppose that $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map and the preference relation $\succsim$ satisfies Axioms 1, 2, 3, 4, 6, and 5. By Theorem 1, we can find $\hat{I} : B(\mathcal{M}, \mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ and $I : B(\mathcal{G}) \times \mathcal{M} \to \mathbb{R}$ satisfying the properties stated in the theorem and such that $f \succsim g$ if and only if $\hat{I}(I(u(f),\cdot)) \geq \hat{I}(I(u(g),\cdot))$ and $f \succsim^m g$ if and all if $I(u(f),m) \geq I(u(g),m)$ for all $f,g \in \mathcal{F}$ and $m \in \mathcal{M}$.

Let $I^m = I(\cdot,m)$ for all $m \in \mathcal{M}$. We show that I^m is translation invariant. Take any $\varphi,\psi \in B(\mathcal{G})$ 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 $\xi \in \{\varphi,\psi\}$, $l \in \{k,r\}$,

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

Then, applying Axiom 6,

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

Then, for any $\varphi',\psi' \in B(\mathcal{G})$ 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:

$$I^m\left(\varphi' + k'\right) = I^m\left(\psi' + k'\right) \implies I^m\left(\varphi' + r'\right) = I^m\left(\psi' + r'\right)\ .$$

Then, take any $\bar{f} \in B(\mathcal{G})$ and $l \in \mathbb{R}$. Normalization implies that $I^m(\bar{f}) = I^m(I^m(\bar{f}))$. By what is shown above, this implies:

$$I^m(\bar{f} + l) = I^m(I^m(\bar{f}) + l) = I^m(\bar{f}) + l$$

proving the claim.

It follows that I^m is normalized, monotone, continuous, quasiconcave, and translation invariant. It is easy to see that Axiom 5 also implies that I^m is monotone continuous. By Theorem 4 in Cerreia-Vioglio et al. (2014), it follows that I^m is a normalized and concave niveloid. Then, by Lemma 26 in Maccheroni et al. (2006), there exists a grounded, lower semicontinuous and convex function $c^m : \Delta \to [0,\infty]$ such that:

$$\begin{aligned}
I^m(\varphi) &= \min_{p' \in \Delta(\mathcal{G})}\left\{\int_\Omega \varphi dp' + c^m(p')\right\}\\
c^m(p) &= \sup_{\varphi' \in B(\mathcal{G})}\left\{I^m(\varphi') - \int_\Omega \varphi' dp\right\}
\end{aligned} \tag{17}$$

for all $\varphi \in B(\mathcal{G})$ and $p \in \Delta(\mathcal{G})$. Then, define $c(\cdot, m) := c^m(\cdot)$ 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 (9).

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

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

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).

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.10: *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 statistical divergence distance.*

Proof of Lemma B.10: 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 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 functions $(\varphi_n)_{n \in \mathbb{N}} \subseteq B_0(\mathcal{G})$ such that for each $n \in \mathbb{N}$, $\varphi_n = -n\chi_E$. Since $m(E) = 0$, $\varphi_n = 0$ a.e. $[m]$ and, therefore, $I^m(\varphi_n) = I^m(0) = 0$ 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,

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

which is a contradiction. 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 Axiom 6 there must exist $f \in \mathcal{F}_0$ such that $fE^m x \succsim x$, 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 $p_0 \ll m$ such 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 $p = m$.

□

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).

(ii) implies (i). That $\succsim$ is a misspecification averse preference follows by Theorem 1. As for monotone continuity, take a sequence of sets $(E_n)_n$ in $\mathcal{G}$ that is monotonically decreasing and such that $E_n \downarrow \emptyset$. Take $f, g \in \mathcal{F}$ such that $f \succ g$ and take any $x \in \mathbb{R}$. Let $\varphi = u(f)$, $\psi = u(g)$, $k = u(x)$, and $\varphi_n = \chi_{E_n} k + \chi_{E_n^c}\varphi$ for all $n \in \mathbb{N}$. Notice that $(\varphi_n)_n$ is uniformly bounded in the norm by $\max\{k, ||\varphi||_\infty\}$ and that $\varphi_n(\omega) \to \varphi(\omega)$ pointwise for all $\omega \in \Omega$. 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 $I(\cdot, m)$ has the Lebesgue property. Therefore: $\xi_n(m) \coloneqq I(\varphi_n, m) \to I(\varphi'_n, m) = \xi'_n(m)$ pointwise for all $m \in \mathcal{M}$. Moreover, monotonicity implies that for all $m \in \mathcal{M}$,

$$\xi_n(m) = I(\varphi_n, m) \geq I(\inf_\omega \varphi_n(\omega), m) = \inf_\omega \varphi_n(\omega) \geq \min\{k, \inf_\omega \varphi(\omega)\} =: L\ .$$

Now, let $\xi' = I(\psi, \cdot)$. Since $f \succ g$, we have that

$$\hat{I}(\xi) = \hat{I}(I(u(f), \cdot) > \hat{I}(I(u(g), \cdot) = \hat{I}(\xi')\ .$$

Since $\hat{I}$ is continuous, we can find $\varepsilon > 0$ small enough such that $\hat{I}(\xi - \varepsilon) > \hat{I}(\xi')$. Set $\xi_\varepsilon = \xi - \varepsilon$. Now, define for each $n \in \mathbb{N}$, the sets

$$\tilde{D}_n = \{m \in \mathcal{M} : |\xi_n(m) - \xi(m)| \geq \varepsilon\}$$

and, then, let $D_n = \cup_{i \geq n} \tilde{D}_i$. Clearly, $(D_n)_n$ is a monotonically decreasing sequence. Moreover, for each given $m \in \mathcal{M}$, since $\xi_i(m) \to \xi(m)$, we can find $N_m \in \mathbb{N}$ such that for all $i \geq N_m$, $|\xi_i(m) - \xi(m)| < \varepsilon$ and, thereby, $m \notin D_n$ for all $n \geq N_m$. It follows that $D_n \downarrow \emptyset$. Since $\hat{I}$ is monotone continuous, we can find $n_0 \in \mathbb{N}$ such that

$$\hat{I}\left(L\chi_{D_{n_0}} + \xi_\varepsilon \chi_{D^c_{n_0}}\right) > \hat{I}(\xi')$$

Now, $\xi_{n_0} \geq L$ by definition and for all $m \in D^c_{n_0}$, since $m \notin \tilde{D}_{n_0}$, it follows that $\xi_n(m) > \xi(m) - \varepsilon = \xi_\varepsilon(m)$. We conclude that $\xi_{n_0}(m) \geq L\chi_{D_{n_0}} + \xi_\varepsilon \chi_{D^c_{n_0}}$ for all $m \in \mathcal{M}$. Since $\hat{I}$ is monotone, we conclude that

$$\hat{I}\left(I\left(u(xE_{n_0}f), \cdot\right)\right) = \hat{I}(\xi_{n_0}) \geq \hat{I}(\xi_\varepsilon) > \hat{I}(\xi') = \hat{I}\left(I(u(g), m)\right) \ .$$

We can, thus, conclude that $xE_{n_0}f$. It follows that $\succsim$ satisfies Axiom 5.

We know want to show that $\succsim$ satisfies variational misspecification. Take now $f, g \in \mathcal{F}$, $x, y \in X$, and $\alpha \in (0,1)$ and assume that $\alpha f + (1-\alpha)x \succsim \alpha g E^m f + (1-\alpha)x$. Then, it follows that

$$\begin{aligned} &\min_{p \ll m} \mathbb{E}_p\left[u(\alpha f + (1-\alpha)x)\right] + c(p,m) \geq \min_{p \ll m} \mathbb{E}_p\left[u(\alpha g + (1-\alpha)x)\right] + c(p,m) \\ &\implies \min_{p \ll m}\{\mathbb{E}_p\left[\alpha u(f)\right] + c(p,m)\} \geq \min_{p \ll m}\{\mathbb{E}_p\left[\alpha u(g)\right] + c(p,m)\} \end{aligned}$$

where we use the fact that u is affine and translation invariance of the variational functional. Then:

$$\begin{aligned} \min_{p \ll m} \mathbb{E}_p\left[u(\alpha f + (1-\alpha)y)\right] + c(p,m) &= \min_{p \ll m} \mathbb{E}_p\left[\alpha u(f)\right] + c(p,m) + (1-\alpha)u(y) \\ &\geq \min_{p \ll m} \mathbb{E}_p\left[\alpha u(g)\right] + c(p,m) + (1-\alpha)u(y) \\ &= \min_{p \ll m} \mathbb{E}_p\left[u(\alpha g + (1-\alpha)y)\right] + c(p,m) \end{aligned}$$

and, therefore, $\alpha f + (1-\alpha)y \succsim \alpha g E^m f + (1-\alpha)y$. This shows that $\succsim$ satisfies $\mathcal{M}$-weak c-independence.

As for the second part, fix $m \in \mathcal{M}$ and suppose there is a $p \in \Delta \setminus \{m\}$ with $p(E^m) = 1$ such that that for all $f \in \mathcal{F}_0$ and $x \in X$, $x \succ \mathbb{E}_p[f]$ implies that $x \succ fE^m x$. In particular, this implies that $\mathbb{E}_p[f] \succsim fE^m\mathbb{E}_p[f]$ for all $f \in \mathcal{F}_0$. Indeed, if this was not the case, we can find $f_0 \in \mathcal{F}_0$ such that $f_0 E^m E_p[f_0] \succ \mathbb{E}_p[f_0]$. Take $x_{f_0,m} \in X$ such that $x_{f_0,m} \sim^m f_0$. Then, $x_{f_0,m} E^m E_p[f_0] \succ \mathbb{E}_p[f_0]$ and, therefore, $u(x_{f_0,m}) > u(E_p[f_0])$.

Thus, $x_{f_0,m} \succ \mathbb{E}_p[f_0]$ and we would obtain $x_{f_0,m} \succ fE^m x_{f_0,m}$, that is, $x_{f_0,m} \succsim^m f$, a contradiction. Then, for all $f \in \mathcal{F}_0$,

$$\mathbb{E}_p[u(f)] = u(\mathbb{E}_p[f]) \geq \min_{q \ll m} \mathbb{E}_q[u(f)] + c(q,m).$$

or, in other words, $\mathbb{E}_p[\varphi] \geq I(\varphi, m)$ for all $\varphi \in B_0(\mathcal{G})$. But, then, using (17),

$$0 \leq c(p,m) = \sup_{\varphi \in B_0(\mathcal{G})} \{I(\varphi,m) - \mathbb{E}_p[\varphi]\} \leq \sup_{\varphi \in B_0(\mathcal{G})} \{\mathbb{E}_p[\varphi] - \mathbb{E}_p[\varphi]\} = 0$$

and, therefore, $c(p,m) = 0$, which is a contradiction. This shows that also part *(ii)* of Axiom 6 is satisfied by $\succsim$. ■

Proof of Corollary 2: The result follows by Proposition 1 and noticing that for each $m \in \mathcal{M}$ and $\varphi \in B(\mathcal{G})$, equation (17) implies that $I_1(\varphi, m) \leq I_2(\varphi, m)$ if and only if $c_1(\cdot, m) \leq c_2(\cdot, m)$. ■

Proof of Proposition 2: Suppose the assumptions of the proposition are satisfied. Pick any $D \in \mathcal{D}_{\mathcal{M}}$. We want to show that there exists a $E^D \in \mathcal{G}$ such that $I^m(\chi_{E^D}) = \chi_D(m)$ for all $m \in \mathcal{M}$. By assumption, there exists $\varphi \in B_0(\mathcal{G})$ such that $0 \leq \varphi \leq 1$ and $I^m(\varphi) = \chi_D(m)$ for all $m \in \mathcal{M}$. Let $E^D = \{\omega \in \Omega : \varphi(\omega) > 0\}$ which clearly is in $\mathcal{G}$. Also notice that $\chi_D \geq \varphi$. Then,using monotonicity, if $m \in D$,

$$1 = \chi_D(m) = I^m(\varphi) \leq I^m(\chi_{E^D}) = \min_p p(E) + c(p,m) \leq m(E^D) \leq 1$$

showing that $\chi_D(m) = 1 = I^m(\chi_{E^D})$. On the other hand, suppose that $m \notin D$. Then

$$0 = \chi_D(m) = I^m(\varphi) = \min_{p \ll m} \int \varphi dp + c(p,m) = \int \varphi d\hat{p} + c(\hat{p}, m)$$

where $\hat{p}$ it the probability where the minimum in the above equation is attained. Then, since $\varphi \geq 0$ and $c(\hat{p}, m) \geq 0$, it must be the case that $\int \varphi d\hat{p} = 0$ which in turn implies that $\hat{p}(E^D) = 0$ and $c(\hat{p}, m) = 0$. Since the latter is uniquely minimized at m, it follows that $m = \hat{p}$ and, therefore, $m(E^D) = 0$. Thus,

$$I^m(\chi_{E^D}) = \min_{p \ll m} p(E^D) + c(p,m) = \min_{p \ll m} 0 + c(p,m) = 0 = \chi_D(m)$$

Since $(\Omega, \mathcal{G})$ is a standard Borel space and $\mathcal{M}$ is a measurable subset of $\mathcal{D}$, then also $(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$ is a standard Borel space[17]. Thus, we can find a sequence $(D_n)_n \subseteq D_{\mathcal{M}}$ that separates points in $\mathcal{M}$. That is, if $m \neq m'$ for some $m, m' \in \mathcal{M}$, there exists D_n such that $m \in D_n$ and $m' \notin D_n$. By defining $\alpha_{\mathcal{M}} : \mathcal{M} \to \{0,1\}^{\mathbb{N}}$, as $\alpha_{\mathcal{M}}(m) = (\chi_{D_n}(m))_n$, the fact that $(D_n)_n$ separates points in $\mathcal{M}$ implies that $\alpha_{\mathcal{M}}$ is injective. It is easy to see that it is also measurable. Since $(\mathcal{M}, \mathcal{D}_{\mathcal{M}})$ is a standard Borel space, the inverse

[17] See Theorems 17.23-17.24 and Corollary 13.4 in Kechris (2012)

$\alpha_{\mathcal{M}}^{-1} : \operatorname{Im}\alpha_{\mathcal{M}} \to \mathcal{M}$ exists and is also measurable. Furthermore, by what is shown above, for each $n \in \mathbb{N}$, we can find $E_n \in \mathcal{G}$ such that $I^m(E_n) = \chi_{D_n}(m)$ for all $m \in \mathcal{M}$. Define similarly $\alpha_\Omega : \Omega \to \{0,1\}^{\mathbb{N}}$ as $\alpha_\Omega(\omega) = (\chi_{E_n}(\omega))_n$, which is also a measurable function. Fix arbitrarily $m_0 \in \mathcal{M}$ and define $\mathfrak{q} : \Omega \to \mathcal{M}$ as

$$\mathfrak{q}(\omega) = \begin{cases} \alpha_{\mathcal{M}}^{-1} \circ \alpha_\Omega(\omega) & \text{if } \alpha_\Omega(\omega) \in \operatorname{Im}\alpha_{\mathcal{M}} \\ m_0 & \text{otherwise.} \end{cases}$$

Clearly, $\mathfrak{q}$ is measurable. We only need to show that $m(\mathfrak{q}^{-1}(m)) = 1$ for all $m \in \mathcal{M}$. Then, fix $m \in \mathcal{M}$. Take $n \in \mathbb{N}$ such that $m \in D_n$. Then, $\chi_{D_n}(m) = 1$ and, therefore,

$$1 \geq m(E_n) = m(E_n) + c(m,m) \geq \min_{p \ll m} p(E_n) + c(p,m) = I^m(\chi_{E_n}) = \chi_{D_n}(m) = 1$$

which implies that $m(E_n) = 1$. On the other hand, take $n \in \mathbb{N}$ such that $m \notin D_n$. Then $\chi_{D_n}(m) = 0$ and, therefore,

$$0 = \chi_{D_n}(m) = I^m(\chi_{E_n}) = \min_{p \ll m} p(E_n) + c(p,m) = \underbrace{\hat{p}(E_n)}_{\geq 0} + \underbrace{c(\hat{p},m)}_{\geq 0}$$

where $\hat{p}$ attains the minimum in the problem $\min_{p \ll m} p(E_n) + c(p,m)$. Then, $\hat{p}(E_n) = 0$ and $c(\hat{p},m) = 0$. But since $c(\cdot,m) \geq 0$ is uniquely minimized at m, it follows that $m = \hat{p}$ and, therefore, $m(E_n) = 0$ and, thereby, $m(\Omega \setminus E_n) = 1$. With this, notice that

$$\begin{aligned}
\mathfrak{q}^{-1}(m) &\supseteq \{\omega \in \Omega : \mathfrak{q}(\omega) = m\} \\
&= \{\omega \in \Omega : \alpha_\Omega(\omega) = \alpha_{\mathcal{M}}(m)\} \\
&= \{\omega \in \Omega : \chi_{E_n}(\omega) = \chi_{D_n}(m) \text{ for all } n \in \mathbb{N}\} \\
&= \bigcap_{n \in \mathbb{N}} \{\omega \in \Omega : \chi_{E_n}(\omega) = \chi_{D_n}(m)\} \\
&= \left(\bigcap_{n : m \in D_n} \{\omega \in \Omega : \chi_{E_n}(\omega) = \chi_{D_n}(m)\} \right) \bigcap \left(\bigcap_{n : m \notin D_n} \{\omega \in \Omega : \chi_{E_n}(\omega) = \chi_{D_n}(m)\} \right) \\
&= (\cap_{n : m \in D_n} E_n) \bigcap (\cap_{n : m \notin D_n} \Omega \setminus E_n)
\end{aligned}$$

and, therefore,

$$1 \geq m\left(\mathfrak{q}^{-1}(m)\right) \geq m\left((\cap_{n : m \in D_n} E_n) \cap (\cap_{n : m \notin D_n} \Omega \setminus E_n)\right) = 1$$

implying the result.

Define $\hat{E}^m = \hat{\mathfrak{q}}^{-1}(m)$ for all m. Fix $m_0 \in \mathcal{M}$ and take $f_1, f_2 \in \mathcal{F}$ and fix any $g \in \mathcal{M}$. Note that since $m_0(\hat{E}^{m_0}) = 1$ and $m(\hat{E}^{m_0}) = 0$ whenever $m \neq m_0$, then

$u(f_i\hat{E}^{m_0}g) = u(f_i)$ a.e.$[m_0]$ and $u(f_i\hat{E}^{m_0}g) = u(g)$ a.e. $[m]$ whenever $m \neq m_0$ for $i = 1, 2$. This implies that $I^{m_0}(u(f_i\hat{E}^{m_0}g)) = I^{m_0}(u(f_i))$ for $i = 1, 2$ and $I^m(u(f_1\hat{E}^{m_0}g)) = I^m(u(f_2\hat{E}^{m_0}g))$ for $m \neq m_0$. Then

$$\begin{aligned} I^{m_0}(u(f_1)) = I^{m_0}(u(f_2)) &\implies I(u(f_1\hat{E}^{m_0}g), \cdot) = I(u(f_2\hat{E}^{m_0}g), \cdot) \\ &\implies \hat{I}\left(I(u(f_1\hat{E}^{m_0}g)\right) = \hat{I}\left(I(u(f_2\hat{E}^{m_0}g)\right) \\ &\implies f_1\hat{E}^{m_0}g \sim f_2\hat{E}^{m_0}g \end{aligned}$$

and

$$\begin{aligned} I^{m_0}(u(f_1)) > I^{m_0}(u(f_2)) &\implies I(u(f_1\hat{E}^{m_0}g), \cdot) > I(u(f_2\hat{E}^{m_0}g), \cdot) \\ &\implies \hat{I}\left(I(u(f_1\hat{E}^{m_0}g)\right) > \hat{I}\left(I(u(f_2\hat{E}^{m_0}g)\right) \\ &\implies f_1\hat{E}^{m_0}g \succ f_2\hat{E}^{m_0}g \end{aligned}$$

This implies that $\forall f_1, f_2, g \in \mathcal{F}$, $f_1\hat{E}^{m_0}g \succsim f_1\hat{E}^m g$ if and only if $I^{m_0}(u(f_1)) \geq I^{m_0}(u(f_2))$ as we wanted to show. Checking the axioms is now routine. ■

Appendix E. Proofs of the Results in Section 5

Proof of Theorem 3: *(i) implies (ii).* Applying Theorem 2, we already know that there exists a convex statistical divergence distance $c : \Delta \times \mathcal{M} \to [0, \infty]$ and a monotone, normalized, and continuous aggregator $\hat{I}$ that satisfies monotone continuity such that $\succsim^m$ is represented by $I(u(f), m) := \min_{p \ll m} \mathbb{E}_p[u(f)] + c(p, m)$ for all $m \in \mathcal{M}$ and $\succsim$ is represented by $\hat{I}(I(u(f), m))$. Define $I(\varphi) = \hat{I}(I(\varphi, m))$. Note that for all $\varphi, \psi \in B(\mathcal{A})$, $I(\varphi, m) \geq I(\psi, m)$ for all $m \in \mathcal{M}$ implies that $I(\varphi) \geq I(\psi)$. Moreover, $I \circ u$ represents the preferences $\succsim$ when restricted to $\mathcal{F}(\mathcal{A})$ and, by Axiom 7, the restriction of $\succsim$ to $\mathcal{F}(\mathcal{A})$ satisfies uncertainty aversion and weak c-independence. A standard proof shows that the restriction of I to $B(\mathcal{A})$ is, then, quasiconcave and translation invariant. By Proposition B.1, it follows that $\hat{I}$ is also quasiconcave and translation invariant. Applying Lemma 26 and Proposition 27 in Maccheroni et al. (2006), we obtain that there exists a grounded, convex, and lower semicontinuous function $\kappa : \Delta(\mathcal{M}) \to [0, \infty]$ such that $\hat{I}(\xi) = \min_{\nu \in \Delta(\mathcal{M})} \int \xi(m) d\nu(m) + \kappa(\nu)$ for all $\xi \in B(\mathcal{M})$. This proves the result.

We already know from Theorem 2 that u is unique up to positive affine transformations and c is unique given u. Uniqueness of κ given u follows by Lemma 26 in Maccheroni et al. (2006).

(ii) implies (i). By Theorem 2, we already know that $\succsim$ is a misspecification averse preferences satisfying Axioms 5 and 6. By Proposition B.1, $I(\varphi) = \hat{I}(I(\varphi, m))$ is

quasiconcave and translation invariant on $B(\mathcal{A})$. Moreover, $I \circ u$ is a representation for the restriction of $\succsim$ to $\mathcal{F}(\mathcal{A})$. A standard argument shows that $\succsim$ satisfies uncertainty aversion and weak c-independence when restricted to $\mathcal{F}(\mathcal{A})$ and, therefore, it satisfies Axiom 7. ∎

Proof of Corollary 3: The result follows from an easy application of Theorem 3, Strzalecki (2011), and standard arguments. ∎

The next Lemma is key in proving Theorem 4.

Lemma B.11: *Suppose $\succsim$ is a misspecification averse preference that satisfies Axioms 5 and 9. There exist an affine and surjective $u : X \to \mathbb{R}$, a strictly increasing $\phi : \mathbb{R} \to \mathbb{R}$, and $\nu \in \Delta(\Omega, \mathcal{A})$ such that for all $f, g \in \mathcal{F}_0(\mathcal{A})$,*

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

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

Proof of Lemma B.11: By Herstein and Milnor (1953), there exists an affine $u : X \to \mathbb{R}$ representing $\succsim$ on X. Since $\succsim$ is unbounded, the argument in the proof of Theorem 1 shows that u must be surjective. Define the binary relation $\succsim^\dagger$ on $B_0(\mathcal{A})$ as follows. For all $\varphi, \psi \in B_0(\mathcal{A})$, $\varphi \succsim^\dagger \psi$ if and only if $f_\varphi \succsim f_\psi$ for some $f_\varphi, f_\psi \in \mathcal{F}_0(\mathcal{A})$ such that $\varphi = u(f_\varphi)$ and $\psi = u(f_\psi)$.

Step 0. $\succsim^\dagger$ is well-defined. That we can find $f_\varphi \in \mathcal{F}_0(\mathcal{A})$ such that $\varphi = u(f)$ for all $\varphi \in B_0(\mathcal{A})$ follows from Lemma B.2. Take $\varphi, \psi \in B_0(\mathcal{A})$ and suppose there are $f_\varphi, f'_\varphi, f_\psi, f'_\psi \in \mathcal{F}$ such that $u(f_\varphi) = \varphi = u(f'_\varphi)$ and $u(f_\psi) = \psi = u(f'_\psi)$. Then, since u represents $\succsim$ on X and by monotonicity,

$$[\forall \omega \in \Omega,\ u(f_\varphi(\omega)) = u(f'_\varphi(\omega))] \implies [\forall \omega \in \Omega,\ f_\varphi(\omega) \sim f'_\varphi(\omega)] \implies f_\varphi \sim f'_\varphi$$
$$[\forall \omega \in \Omega,\ u(f_\psi(\omega)) = u(f'_\psi(\omega))] \implies [\forall \omega \in \Omega,\ f_\psi(\omega) \sim f'_\psi(\omega)] \implies f_\psi \sim f'_\psi$$

We conclude that $f_\varphi \succsim f_\psi$ if and only if $f'_\varphi \succsim f'_\psi$, showing that $\succsim^\dagger$ is well-defined.

Step 1. $\succsim^\dagger$ is complete and transitive. Take $\varphi, \psi \in B_0(\mathcal{A})$. By Lemma B.2, we can find $f_\varphi, f_\psi \in \mathcal{F}_0(\mathcal{A})$ such that $\varphi = u(f_\varphi)$ and $\psi = u(f_\psi)$. Since $\succsim$ is complete, either $f_\varphi \succsim f_\psi$, and then $\varphi \succsim^\dagger \psi$, or $f_\psi \succsim f_\varphi$, and then $\psi \succsim^\dagger \varphi$. This shows that $\succsim^\dagger$ is complete. As for transitivity, take $\varphi_1, \varphi_2, \varphi_3 \in B_0(\mathcal{A})$ such that $\varphi_1 \succsim^\dagger \varphi_2$ and $\varphi_2 \succsim^\dagger \varphi_3$. Then, there exist $f_i \in \mathcal{F}_0(\mathcal{A})$ such that $\varphi_i = u(f_i)$ for each $i \in \{1, 2, 3\}$ and $f_1 \succsim f_2$ and $f_2 \succsim f_3$. Since $\succsim$ is transitive, $f_1 \succsim f_3$ and, therefore, $\varphi_1 \succsim^\dagger \varphi_3$.

Step 2. For all $E \in \mathcal{A}$, E is $\succsim$-null if and only if E is $\succsim^\dagger$-null. Suppose that E is $\succsim$-null. Take any $\varphi, \psi \in B_0(\mathcal{A})$. Then, we can find f_φ, f_ψ such that $\varphi = u(f_\varphi)$ and

$\psi = u(f_\psi)$. Since E is $\succsim$-null, $f_\varphi E f_\psi \sim f_\psi$. Finally, note that $u(f_\varphi E f_\psi) = \varphi E \psi$ and, therefore, $varphiE\psi \sim^\dagger \psi$. On the other hand, suppose that $E \in \mathcal{A}$ is $\succsim^\dagger$-null. Then, for any $fmg \in \mathcal{F}_0(\mathcal{A})$, we have that $u(fEg) = u(f)Eu(g) \sim^\dagger u(g)$. It follows that $fEg \sim g$.

Step 3. Tradeoff Consistency Take $A, E \in \mathcal{A}$ that are nonnull for $\succsim^\dagger$, $r_1, r_2, t_1, t_2 \in \mathbb{R}$, $\varphi, \psi \in B_0(\mathcal{A})$, and assume that $r_1 A\varphi \succsim^\dagger t_1 A\psi$, $r_2 Af\varphi \succsim^\dagger t_2 A\psi$, and $r_1 E\varphi \succsim^\dagger t_1 E\psi$. Take now $x_1, x_2, y_1, y_2 \in X$ such that $r_i = u(x_i)$ and $t_i = u(y_i)$ for $i \in \{1, 2\}$ and f_φ, f_ψ such that $\varphi = u(f_\varphi)$ and $\psi = u(f_\psi)$. It follows that $x_1 Af_\varphi \succsim^\dagger y_1 Af_\psi$, $x_2 Af_\varphi \succsim^\dagger y_2 Af_\psi$, and $x_1 Ef_\varphi \succsim^\dagger y_1 Ef_\psi$. Moreover, by Step 2, we also know that E and A are nonnull for $\succsim$. By Axiom, we conclude that $x_2 Ef_\varphi \succsim^\dagger y_2 Ef_\psi$ and, therefore, $r_2 E\varphi \succsim^\dagger t_2 E\psi$. This proves the step.

Step 3. S-Continuity Fix a finite, measurable partition $\{E_1, \dots E_n\}$ and take $\varphi = \sum_{i=1}^n \chi_{E_i} r_i \in B_0(\mathcal{A})$, with $r_i \in \mathbb{R}$ for all $i \in \{1, \dots, n\}$. We want to show that

$$A \coloneqq \{(t_1, \dots t_n) \in \mathbb{R}^n : \sum_{i=1}^n \chi_{E_i} t_i \succsim^\dagger \varphi\} \quad \text{and} \quad B \coloneqq \{(t_1, \dots t_n) \in \mathbb{R}^n : \varphi \succsim^\dagger \sum_{i=1}^n \chi_{E_i} t_i\}$$

are closed in $\mathbb{R}^n$. First of all, since u is surjective, we can pick $z_1, \dots, z_n \in X$ such that $r_i = u(z_i)$ for all $i \in \{1, \dots, n\}$. Then, if we let $f = \sum_{i=1}^n \chi_i z_i$, we see that $f \in \mathcal{F}$ and $\varphi = u(f)$. By Lemma, we can find $z \in X$ such that $f \sim z$ and, therefore, $\varphi \sim^\dagger r_\varphi \coloneqq u(z)$. To show that A is closed, take a sequence $(t^k)_k$ in $\mathbb{R}^n$ such that $t^k \to t \in \mathbb{R}^n$ and $\sum_{i=1}^n \chi_{E_i} t_i^k \succsim^\dagger \varphi$ for all $k \in \mathbb{N}$. We want to show that $\sum_{i=1}^n \chi_{E_i} t_i \succsim^\dagger \varphi$. Since $t^k \to t$ in the product topology, the sequence $(t^k)_k$ is supnorm bounded and, therefore, we can find $l, L \in \mathbb{R}$ such that $l \leq t_i^k \leq L$ for all $i \in \{1, \dots n\}$ and $k \in \mathbb{N}$. In particular, for each $i \in \{1, \dots n\}$ and $k \in \mathbb{N}$, we can find $\alpha_i^k \in [0, 1]$ such that $t_i^k = \alpha_i^k L + (1 - \alpha_i^k) l$. Similarly, we can find $\alpha \in [0, 1]^n$ such that $t_i = \alpha_i L + (1 - \alpha_i) l$ for all $i \in \{1, \dots, n\}$. It is easy to see that $\alpha^k \to \alpha$. Moreover, Since u is surjective, we can find $x, y \in X$ such that $u(x) = L$ and $u(y) = l$ and, since u is affine, for all $i \in \{1, \dots n\}$ it holds that

$$u(\alpha_i^k x + (1 - \alpha_i^k) y) = \alpha_i^k u(x) + (1 - \alpha_i^k) u(y) = \alpha_i^k L + (1 - \alpha_i^k) l = t_i^k$$

for all $k \in \mathbb{N}$ and, similarly, that $t_i = \alpha_i u(x) + (1-\alpha_i)u(y)$. Then:

$$\begin{aligned}\sum_{i=1}^n \chi_{E_i} t_i^k \succsim^\dagger \varphi &\iff \sum_{i=1}^n \chi_{E_i} t_i^k \succsim^\dagger r_\varphi \\ &\iff \sum_{i=1}^n \chi_{E_i} u(\alpha_i^k x + (1-\alpha_i^k)y) \succsim^\dagger u(z) \\ &\iff \sum_{i=1}^n \chi_{E_i} [\alpha_i^k x + (1-\alpha_i^k)y] \succsim z\end{aligned}$$

for all $k \in \mathbb{N}$. By S-continuity, we conclude that $\sum_{i=1}^n \chi_{E_i}[\alpha_i x + (1-\alpha_i)y] \succsim z$ and, therefore, $\sum_{i=1}^n \chi_{E_i} t_i \succsim^\dagger r_\varphi \sim^\dagger \varphi$. It follows that A is closed. An analogous argument shows that also B is closed.

Clearly, $\mathbb{R}$ is connected and separable. Moreover, $\succsim^\dagger$ is a weak order on $B_0(\mathcal{A})$ that satisfies S-Continuity and Tradeoff Consistency. By Theorem V.3.4 in Wakker, there exists a finitely additive probability ν and a continuous function $\phi : \mathbb{R} \to \mathbb{R}$ such that $J(\varphi) = \int \phi(\varphi) d\nu$ represents $\succsim^\dagger$ on $B_0(\mathcal{A})$. Moreover, by assumption, there are two $\succsim$-nonnull, disjoint events in $\mathcal{A}$. By the previous step, these two events are also $\succsim^\dagger$-nonnull. Then, Observation V.3.4' in Wakker implies that ν is unique and ϕ is cardinally unique. Monotone continuity implies that ν is countably additive.

In particular, for all $f, g \in \mathcal{F}_0(\mathcal{A})$,

$$f \succsim g \iff u(f) \succsim^\dagger u(g) \iff \int \phi(u(f(\omega)))d\nu(\omega) \geq \int \phi(u(g(\omega)))d\nu(\omega)$$

It is evident that ϕ is strictly increasing. ■

Lemma B.12: *Suppose $(\Omega, \mathcal{G}, \mathcal{M})$ admits a best-fit map and there exist a utility function $u : X \to \mathbb{R}$, a convex statistical divergence 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 $\mu \in \Delta(\mathcal{M}, \mathcal{D}_\mathcal{M})$ 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 (9). Then, there exists a probability measure $\nu \in \Delta(\Omega, \mathcal{A})$ 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.12: 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 $\nu \in \Delta(\Omega, \mathcal{A})$ 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 Lemma B.6. 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 4: *(i) implies (ii).*

We know that $\succsim$ is represented by u when restricted to constant acts. Define the functional $I : B(\mathcal{G}) \to \mathbb{R}$ such that for each $\varphi \in B(\mathcal{G})$, $I(\varphi) \coloneqq u(x_{f_\varphi})$, 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 $\varphi \in B(\mathcal{G})$, while the certainty equivalent $x_{f_\varphi} \sim f_\varphi$ exists by Lemma B.3. Moreover, for any $\varphi \in B(\mathcal{G})$, 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

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

showing that I is a well-defined functional on $B(\mathcal{G})$. It is easily seen that such functional is also normalized, monotone, and continuous.[18]

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 Theorem 2, 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 (17). Moreover, let $I_{\mathcal{A}}$ be the generalized conditional expectation from Lemma B.6. Take now $\varphi, \psi \in B(\mathcal{G})$ 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

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

$f_\varphi \succsim^m f_\psi$ 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.11, there exist an unbounded and affine $\tilde{u} : X \to \mathbb{R}$, a strictly increasing $\phi : \mathbb{R} \to \mathbb{R}$, and a probability $\nu \in \Delta(\Omega, \mathcal{A})$ such that the restriction of $\succsim$ to $\mathcal{F}_0(\mathcal{A})$ 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 B.6, 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 $\nu(E) = \mu(D)$.

Uniqueness follows by standard arguments.

(ii) implies (i). It is clear that $\succsim$ satisfies Axioms 1-5 are satisfied. Moreover, there exists a probability measure $\nu \in \Delta(\Omega, \mathcal{A})$ 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 Wakker's axioms when restricted to $\mathcal{F}(\mathcal{A})$. ■

Proof of Theorem 5: *(i) implies (ii).* We know that there exists an affine $u : X \to \mathbb{R}$ and a normalized, monotone, continuous, functional $I : B(\mathcal{G}) \to \mathbb{R}$ such that $\succsim$ is represented by $I \circ u$ on $\mathcal{F}$. By Theorem 2, we know that for each $m \in \mathcal{M}$, there exists I^m given as in (9) such that $I^m \circ u$ represents $\succsim^m$ on $\mathcal{F}$. By consistency, we also know that for all $\varphi, \psi \in B(\mathcal{G})$, $I^m(\varphi) \geq I^m(\psi)$ for all $m \in \mathcal{M}$ implies that $I(\varphi) \geq I(\psi)$. Therefore, by Proposition B.1, there exists a unique normalized, monotone, and continuous $\hat{I} : B(\mathcal{D}_{\mathcal{M}}) \to \mathbb{R}$ such that $\hat{I}(I(\varphi, \cdot)) = I(\varphi)$ for all $\varphi \in B(\mathcal{A})$. Take $\xi \in B(\mathcal{D}_{\mathcal{M}})$. By Lemma B.9, we can find a $\varphi \in B(\mathcal{A})$ such that $\xi = T(\varphi)$ and $f \in \mathcal{F}$ 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 $\hat{I}(\xi) = r_0 = \inf_{m \in \mathcal{M}} \xi(m)$. 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 $\varphi \in B(\mathcal{G})$,

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

Suppose that in addition $\mathcal{M}$ is compact. Since c is jointly lower semicontinuous, applying 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:

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

Since $\varphi \in B(\mathcal{G})$ was arbitrarily chosen, we conclude that this holds everywhere on $B(\mathcal{G})$. 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.

(ii) implies (i). It follows by routine arguments. ■

Appendix F. Proofs of the Results in Section 6

Proof of Proposition 3: For all $p \geq 0$, define the function

$$\Phi_p(z, \alpha) = pz + \lambda\eta \left((1 - z) \log \frac{(1 - z)}{\alpha} + z \log \frac{z}{1 - \alpha} \right) + \lambda (\gamma - p) \left[\frac{1}{2} - z \right]^+$$

for all $z \in [0, 1]$ and $\alpha \in [0, 1/2]$.

Pick any $0 < p \leq \gamma$ and consider $m \in \mathcal{M}$. Let $A = [0, p)$. Since γ is a median of m, we have that $m([\gamma, \infty)) \geq 1/2$ and, therefore,

$$m(A) = 1 - m([p, \infty)) \leq 1 - m([\gamma, \infty)) \leq \frac{1}{2} \, .$$

Put $\alpha_m = m(A)$ and note that $\alpha_m \in [0, 1/2]$. Fix $q \ll m$ and put $z_q = q(A^c)$. By the data-processing inequality, we have

$$\begin{aligned} R(q || m) &\geq R(q|_{\sigma(\{A\})} || m_{\sigma(\{A\})}) \\ &= q(A) \log \frac{q(A)}{m(A)} + q(A^c) \log \frac{q(A^c)}{m(A^c)} \\ &= (1 - z_q) \log \frac{1 - z_q}{\alpha_m} + z_q \log \frac{z_q}{1 - \alpha_m} \, . \end{aligned}$$

On the other hand,

$$\begin{aligned} W_1(q, m) &= \inf_{\pi \in \Pi(q,m)} \int |\omega_1 - \omega_2| d\pi(\omega_1, \omega_2) \\ &\geq \inf_{\pi \in \Pi(q,m)} \int_{\omega_1 \leq p, \omega_2 \geq \gamma} |\omega_1 - \omega_2| d\pi(\omega_1, \omega_2) \\ &\geq \inf_{\pi \in \Pi(q,m)} \int_{\omega_1 \leq p, \omega_2 \geq \gamma} (\gamma - p) d\pi(\omega_1, \omega_2) \\ &= (\gamma - p) \inf_{\pi \in \Pi(q,m)} \pi(\{\omega_1 \leq p, \omega_2 \geq \gamma\}) \end{aligned}$$

where recall that $\gamma - p \geq 0$. Moreover, for any $\pi \in \Pi(q, m)$, we have

$$\begin{aligned}\pi(\{\omega_1 \leq p, \omega_2 \geq \gamma\}) &= \pi\left(\{\omega_2 \geq \gamma\}\right) - \pi\left(\{\omega_2 \geq \gamma\} \cap \{\omega_1 > p\}\right) \\ &\geq \pi\left(\{\omega_2 \geq \gamma\}\right) - \pi\left(\{\omega_1 \geq p\}\right) \\ &= m([\gamma, \infty)) - q([p, \infty)) \\ &\geq \frac{1}{2} - z_q\end{aligned}$$

and, therefore,

$$W_1(q, m) \geq (\gamma - p)\left[\frac{1}{2} - z_q\right]^+ .$$

Putting the pieces together, we have that:

$$\begin{aligned}pq(A_p^c) + \lambda c_\eta(q, m) &\geq pq(A_p^c) + \lambda\eta\left(q(A_p)\log\frac{q(A_p)}{m(A_p)} + q(A_p^c)\log\frac{q(A_p^c)}{m(A_p^c)}\right) + \lambda(\gamma - p)\left[\frac{1}{2} - q(A_p^c)\right]^+ \\ &\geq \inf_{\alpha \in [0, \frac{1}{2}]} \inf_{z \in [0,1]} \Phi(z, \alpha) .\end{aligned}$$

Let $\hat{\Phi}_p(\alpha) = \min_{z \in [0,1]} \Phi_p(z, \alpha)$. For each given $\alpha \in (0, \frac{1}{2})$, $\Phi_p(z, \alpha)$ is convex in z and strictly convex on both $(0, \frac{1}{2})$ and on $(\frac{1}{2}, 1)$. It follows that the problem admits a unique minimizer $\hat{z}(\alpha)$ for all $\alpha \in (0, \frac{1}{2})$ which is either equal to $\frac{1}{2}$ or interior in either $(0, \frac{1}{2}$ or in $(\frac{1}{2}, 1)$. Now, consider $\alpha \in (0, \alpha)$. If $\hat{z}(\alpha) \in (0, \frac{1}{2}]$, then $\hat{z}(\alpha) \leq \frac{1}{2} < 1 - \alpha$. If, instead, $\hat{z}(\alpha) \in (\frac{1}{2}, 1)$, it needs to satisfy the FOC:

$$\begin{aligned}&\frac{\partial \Phi_p(\hat{z}(\alpha), \alpha)}{\partial z} = p + \lambda\eta\left[\log\frac{\hat{z}(\alpha)}{1 - \hat{z}(\alpha)} + \log\frac{\alpha}{1 - \alpha}\right] = 0 \\ &\implies \frac{\hat{z}(\alpha)}{1 - \hat{z}(\alpha)} = \exp\left\{-\frac{p}{\lambda\eta} - \log\frac{\alpha}{1 - \alpha}\right\} = e^{-\frac{p}{\lambda\eta}}\,\frac{1 - \alpha}{\alpha} < \frac{1 - \alpha}{\alpha} \\ &\implies \hat{z}(\alpha) < 1 - \alpha .\end{aligned}$$

We conclude that $\hat{z}(\alpha) < 1 - \alpha$ for all $\alpha \in (0, \frac{1}{2})$. Now, for $\alpha \in (0, \frac{1}{2})$, applying the envelope theorem we obtain that:

$$\frac{d\hat{\Phi}_p(\alpha)}{d\alpha} = \frac{\partial \Phi_p(\hat{z}(\alpha), \alpha)}{\partial \alpha} = \frac{\hat{z}(\alpha) - (1 - \alpha)}{\alpha(1 - \alpha)} < 0 .$$

Moreover, for $\alpha = 0$, we have that $\Phi_p(z, 0) = +\infty$ if $z \neq 1$ and $\Phi_p(z, 0) = p$ if $z = 1$. It follows that $\hat{z}(\alpha) = 1$. Moreover, for all $\alpha \in (0, \frac{1}{2})$,

$$0 \leq \hat{\Phi}_p(\alpha) \leq \Phi_p(1, \alpha) = p - \log(1 - \alpha)$$

and, therefore, taking limits $\lim_{\alpha\downarrow 0}\hat{\Phi}_p(\alpha)\leq p$. We conclude that

$$\inf_{\alpha\in[0,\frac{1}{2}]}\inf_{z\in[0,1]}\Phi_p(z,\alpha)=\inf_{\alpha\in[0,\frac{1}{2}]}\hat{\Phi}(\alpha)=\hat{\Phi}\left(\frac{1}{2}\right)=\min_{z\in[0,1]}\Phi_p\left(z,\frac{1}{2}\right)=\min_{z\in[0,\frac{1}{2}]}\Phi_p\left(z,\frac{1}{2}\right)$$

where the last equality follows from the fact that we showed above that $\hat{z}(\frac{1}{2})\leq 1-\frac{1}{2}=\frac{1}{2}$. We conclude that

$$pq(A_p^c)+\lambda c_\eta(q,m)\geq\min_{z\in[0,\frac{1}{2}]}\Phi_p\left(z,\frac{1}{2}\right)\ .$$

Since m and $q\ll m$ were chosen arbitrarily, this holds for all $m\in\mathcal{M}$ and $q\ll m$ and, therefore, we conclude that

$$V(f_p)=\inf_{m\in\mathcal{M}}\inf_{q\ll m}pq(A_p^c)+\lambda c_\eta(q,m)\geq\min_{z\in[0,\frac{1}{2}]}\Phi_p\left(z,\frac{1}{2}\right)\ .$$

Moreover, fix any $z\in[0,\frac{1}{2}]$ and take $m_\varepsilon=\frac{1}{2}\delta_{p-\varepsilon}+\frac{1}{2}\delta_\gamma$ and $q_\varepsilon=(1-z)\delta_{p-\varepsilon}+z\delta_\gamma$ for $0<\varepsilon<p$. Then:

$$R(q_\varepsilon,m_\varepsilon)=(1-z)\log\frac{1-z}{1/2}+z\log\frac{z}{1/2}$$

and

$$W_1(q_\varepsilon,m_\varepsilon)=\int|F_{q_\varepsilon}(t)-F_{m_\varepsilon}(t)|dt=\left(\frac{1}{2}-z\right)(\gamma-p+\varepsilon)$$

so that

$$V(f_p)\leq pz+\lambda\left[\left(\frac{1}{2}-z\right)(\gamma-p+\varepsilon)\right]+\lambda\eta\left[(1-z)\log\frac{1-z}{1/2}+z\log\frac{z}{1/2}\right]\ .$$

Taking limits as $\varepsilon\downarrow 0$, we ge that

$$V(f_p)\leq\Phi_p(z,\alpha)\leq\min_{z\in[0,\frac{1}{2}]}\Phi_p\left(z,\frac{1}{2}\right)$$

showing that, indeed,

$$V(f_p)=\min_{z\in[0,\frac{1}{2}]}\Phi_p\left(z,\frac{1}{2}\right)$$

where recall that

$$\begin{aligned}\Phi_p\left(z,\frac{1}{2}\right)&=pz+\lambda\eta\left[(1-z)\log 2(1-z)+z\log 2z\right]+\lambda(\gamma-p)\left(\frac{1}{2}-z\right)\\&=\left[(1+\lambda)p-\lambda\gamma\right]z+\lambda\eta\left[(1-z)\log(1-z)+z\log z\right]+\lambda\eta\log 2+\lambda\frac{\gamma-p}{2}.\end{aligned}$$

Now,

$$\frac{\partial\Phi_p(z,1/2)}{\partial z}=(1+\lambda)p-\lambda\gamma+\lambda\eta\log\frac{z}{1-z}$$

and, therefore,

$$\frac{\partial \Phi_p(z,1/2)}{\partial z} < 0 \iff z < \frac{1}{1+\exp\{\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}\}} \; .$$

Notice that the RHS is above 1/2 if and only if

$$\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta} \leq 0 \iff p \leq \frac{\lambda}{1+\lambda}\gamma \; .$$

Thus, we have two cases:

$$\begin{aligned} p \leq \frac{\lambda}{1+\lambda}\gamma &\implies \hat{z} = \frac{1}{2} \; , \\ p > \frac{\lambda}{1+\lambda}\gamma &\implies \hat{z} = \frac{1}{1+\exp\{\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}\}} \; . \end{aligned}$$

For the second case, since FOC must hold, we have that $(1+\lambda)p - \lambda\gamma = \lambda\eta \log\frac{1-\hat{z}}{\hat{z}}$ and, therefore,

$$\begin{aligned} [(1+\lambda)p-\lambda\gamma]\,\hat{z} + \lambda\eta\,[(1-\hat{z})\log(1-\hat{z}) + \hat{z}\log\hat{z}] &= \lambda\eta\left[\hat{z}\log\frac{1-\hat{z}}{\hat{z}} + (1-\hat{z})\log(1-\hat{z}) + \hat{z}\log\hat{z}\right] \\ &= \lambda\eta\log(1-\hat{z}) \\ &= \lambda\eta\log\frac{1}{1+\exp\{-\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}\}} \\ &= -\lambda\eta\log\left(1+e^{-\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}}\right) \end{aligned}$$

and, therefore,

$$V(f_p) = \begin{cases} \frac{p}{2} & \text{if } 0 < p \leq \frac{\lambda}{1+\lambda}\gamma \\ \lambda\eta\left[\log 2 - \log\left(1+e^{-\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}}\right)\right] + \lambda\frac{\gamma-p}{2} & \text{if } \frac{\lambda}{1+\lambda}\gamma < p \leq \gamma \end{cases} \; .$$

Note first that for any $p \geq 0$, we have that for any $m \in \Updownarrow$ and $q \ll m$,

$$pq(\omega \geq p) + c_\eta(q,m) \geq 0$$

so that $V(f_p) \geq 0$. If $p = 0$, choose for example $m_\gamma = q_\gamma = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_\gamma$ and notice that

$$0q_\gamma(\omega \geq p) + c_\eta(q_\gamma, m_\gamma) = 0$$

so that we can conclude that $V(f_0) = 0$. Similarly, for $p > \gamma$,

$$pq(\omega \geq q_\gamma) + c_\eta(q_\gamma, m_\gamma) = 0$$

and, therefore, $V(f_p) = 0$ for all $p > \gamma$.

We conclude that

$$V(f_p) = \begin{cases} \frac{p}{2} & \text{if } 0 \leq p \leq \frac{\lambda}{1+\lambda}\gamma \\ \lambda\eta\left[\log 2 - \log\left(1 + e^{-\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}}\right)\right] + \lambda\frac{\gamma-p}{2} & \text{if } \frac{\lambda}{1+\lambda}\gamma < p \leq \gamma \\ 0 & \text{if } p > \gamma \end{cases} .$$

We can, therefore, restrict attention to $p \in [0, \gamma]$. First of all, take the derivative for $\frac{\lambda}{1+\lambda}\gamma < p < \gamma$:

$$\frac{\partial V(f_p)}{\partial p} = \frac{1+\lambda}{1 + e^{\frac{(1+\lambda)p-\lambda\gamma}{\lambda\eta}}} - \frac{\lambda}{2}$$

which is positive if and only if

$$p \leq \frac{\lambda\gamma + \lambda\eta\log(1 + 2/\lambda)}{1+\lambda} .$$

Notice that

$$\frac{\lambda\gamma + \lambda\eta\log(1 + 2/\lambda)}{1+\lambda} < \gamma \iff \lambda\eta\log(1 + 2/\lambda) < \gamma .$$

Moreover, since $\lambda\log(1 + 2/\lambda) < 2$ for all $\lambda > 0$ and $\eta < \gamma/2$, it is, indeed, the case that $\lambda\eta\log(1 + 2/\lambda) < \gamma$ and, therefore, conclude that the solution is

$$\hat{p}(\lambda) = \frac{\lambda\gamma + \lambda\eta\log(1 + 2/\lambda)}{1+\lambda} \in (0, \gamma) .$$

Moreover:

$$\begin{aligned} \frac{d\hat{p}(\lambda)}{d\lambda} &= \frac{\left(\gamma + \eta\log(1 + 2/\lambda) - \frac{2\eta}{2+\lambda}\right)(1+\lambda) - (\lambda\gamma + \lambda\eta\log(1 + 2/\lambda))}{(1+\lambda)^2} \\ &= \frac{\gamma + \eta\log(1 + 2/\lambda) - \eta\frac{2+2\lambda}{2+\lambda}}{(1+\lambda)^2} \end{aligned}$$

Let $\psi(\lambda) = \gamma + \eta\log(1 + 2/\lambda) - \eta\frac{2+2\lambda}{2+\lambda}$ and notice that

$$\frac{d\psi(\lambda)}{d\lambda} = -\frac{2\eta}{\lambda + \lambda^2} - \frac{2\eta}{(2+\lambda)^2} < 0$$

for all $\lambda > 0$ and that

$$\lim_{\lambda\to 0} \psi(\lambda) = \infty \quad \text{and} \quad \lim_{\lambda\to\infty} \psi(\lambda) = \gamma - 2\eta > 0 .$$

It follows that $\frac{d\hat{p}(\lambda)}{d\lambda} > 0$ for all $\lambda > 0$. Finally,

$$\lim_{\lambda\to 0} \hat{p}(\lambda) = 0 \quad \text{and} \quad \lim_{\lambda\to\infty} \hat{p}(\lambda) = \gamma.$$

■