Making Decisions under Model Misspecification*

Simone Cerreia–Vioglio[a], Lars Peter Hansen[b], Fabio Maccheroni[a]
and Massimo Marinacci[a]

[a]Università Bocconi and Igier, [b]University of Chicago

March 2025

Abstract

We use decision theory to confront uncertainty that is sufficiently broad to incorporate "models as approximations." We presume the existence of a featured collection of what we call "structured models" that have explicit substantive motivations. The decision maker confronts uncertainty through the lens of these models, but also views these models as simplifications, and hence, as misspecified. We extend the max-min analysis under model ambiguity to incorporate the uncertainty induced by acknowledging that the models used in decision-making are simplified approximations. Formally, we provide an axiomatic rationale for a decision criterion that incorporates model misspecification concerns. We then extend our analysis beyond the max-min case allowing for a more general criterion that encompasses a Bayesian formulation.

Keywords: *uncertainty, decision theory, model misspecification, ambiguity.*

JEL codes— C54, D81

First Draft: December 2019. We thank Attila Ambrus, Pierpaolo Battigalli, Benjamin Brooks, Tim Christensen, Roberto Corrao, Giacomo Lanzani, Marco Loseto, Philipp Sadowski, Todd Sarver, Jesse Shapiro as well as the audiences at Advances in Decision Analysis 2019, SAET 2019, Blue Collar Working Group 2.0, One World Mathematical Game Theory Seminar, RUD 2020, UCL-Osaka International Conference on the Mathematics for Risk and Decisions, MUSEES 2022, Confronting Uncertainty in Climate Change, Bicocca, Bilkent, Caltech, Duke, Georgetown, Glasgow, Harvard, Johns Hopkins, LSE, MIT, Oxford, Parma, UAB, UCL, Warwick for their very useful comments. We thank for the financial support the Alfred P. Sloan Foundation (grant G-2018-11113), the European Research Council (grants SDDM-TEA and INDIMACRO) and the Ministero dell'università e della ricerca (grant 2017CY2NCA).

Come l'araba fenice:
che vi sia, ciascun lo dice;
dove sia, nessun lo sa.[1]

1 Introduction

The consequences of a decision may depend on exogenous contingencies and uncertain outcomes that are outside the control of a decision maker. This uncertainty takes on many forms. Economic applications typically feature *risk*, where the decision maker knows the correct probabilistic model governing the contingencies but not necessarily the decision outcomes. Yet, this is a demanding assumption. As a result, statisticians and econometricians have long wrestled with how to confront *ambiguity* over models or unknown parameters within a model. Each model is itself a simplification or an approximation designed to guide or enhance our understanding of some underlying phenomenon of interest. Thus, the model, by its very nature, is *misspecified*, but in typically uncertain ways. How should a decision maker acknowledge model misspecification in a way that guides the use of purposefully simplified models sensibly? This concern has certainly been on the radar screen of statisticians and control theorists, but it has been largely absent in formal approaches to decision theory.[2] Indeed, the statisticians Box and Cox have both stated the challenge succinctly in complementary ways:

> Since all models are wrong, the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad. Box (1976).
>
> ... it does not seem helpful just to say that all models are wrong. The very word "model" implies simplification and idealization. The idea that complex physical, biological or sociological systems can be exactly described by a few formulae is patently absurd. The construction of idealized representations that capture important stable aspects of such systems is, however, a vital part of general scientific analysis and statistical models, especially substantive ones ... Cox (1995).

While there are formulations of decision and control problems that intend to confront model misspecification, the aim of this paper is: (i) to develop an axiomatic approach that will provide a rigorous guide for applications and (ii) to enrich formal decision theory when applied to environments with uncertainty through the guise of models.

[1] "Like the Arabian phoenix: that it exists, everyone says; where it is, nobody knows." A passage from a libretto of Pietro Metastasio.

[2] In Hansen (2014) and Hansen and Marinacci (2016) three kinds of uncertainty are distinguished based on the knowledge of the decision maker, the most challenging being model misspecification viewed as uncertainty induced by the approximate nature of the models under consideration.

The protagonist of our analysis is a decision maker who is able to formulate models – for instance a policy maker having to decide a climate policy based on existing alternative climate models – but is concerned about their misspecification and wants to use a decision criterion which accounts for that. Our axiomatic analysis, which has a normative nature, aims to derive a criterion of this kind to help the decision maker to cope with model misspecification in a principled way. In this endeavour, we follow Hansen and Sargent (2022) by referring to the formulated models as "structured models." These structured models are ones that are explicitly motivated or featured, such as ones with substantive motivation or scientific underpinnings, consistent with the use of the term "models" by Box and Cox. They may be based on scientific knowledge relying on empirical evidence and theoretical arguments or on revealing parameterizations of probability models with parameters that are interpretable to the decision maker. In posing decision problems formally, it is often assumed, following Wald (1950), that the correct model belongs to the set of models that decision makers posit. The presumption that a decision maker identifies, among their hypotheses, the correct model is often questionable – recalling the initial quotation, the correct model is often a decision maker phoenix. We embrace, rather than push aside, the "models are approximations" perspective of many applied researchers, as articulated by Box, Cox and others. To explore misspecification formally, we introduce a potentially rich collection of probability distributions that depict possible representations of the data without formal substantive motivation. We refer to these as "unstructured models." We use such alternative models as a way to capture how models could be misspecified.[3]

This distinction between structured and unstructured is central to the analysis in this paper and is used to distinguish aversion to ambiguity over models and aversion to potential model misspecification. At a decision-theoretic level, a proper analysis of misspecification concerns has remained elusive so far. Indeed, many studies dealing with economic agents confronting model misspecification still assume that they are conventional expected utility decision makers who do not address formally potential model misspecification concerns in their preference ordering.[4] We extend the analysis of Hansen and Sargent (2022) by providing an axiomatic underpinning for a corresponding decision theory along with a representation of the implied preferences that can guide applications. In so doing, we show an important connection with the analysis of subjective and objective rationality of Gilboa et al. (2010).

Criterion This paper proposes a first decision-theoretic analysis of decision making under model misspecification. We consider a classical setup in the spirit of Wald (1950), but relative to his seminal work we explicitly remove the assumption that the correct model belongs to

[3]Such a distinction is also present in earlier work by Hansen and Sargent (2007) and Hansen and Miao (2018) but without specific reference to the terms "structured" and "unstructured."

[4]See, e.g., Esponda and Pouzo (2016) and Fudenberg et al. (2017). In contrast, using a decision criterion elaborated in this paper (cf. Section 6.1), Lanzani (2023) has recently showed the relevance for these analyses of a proper account of misspecification concerns.

the set of posited models and we allow for nonneutrality toward this feature. More formally, in our purely normative approach we assume that decision makers posit a set Q of *structured* (probabilistic) *models* q on states, motivated by their information, but they are afraid that none of them is correct and so face model misspecification. For this reason, decision makers contemplate what we call *unstructured models* in ranking acts f, according to a conservative decision criterion

$$V(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \min_{q \in Q} c(p, q) \right\} \tag{1}$$

where Δ denotes the set of all probabilities. To interpret this criterion, let

$$C(p, Q) = \min_{q \in Q} c(p, q)$$

where we presume that $C(q, Q) = 0$ if and only if $q \in Q$. In this construction, $C(p, Q)$ is what we call a Hausdorff statistical set distance between a model p and the posited set Q of structured models. This distance is nonzero if and only if p is unstructured, that is, $p \notin Q$. More generally, p's that are closer to the set of structured models Q have a less adverse impact on the preferences, as is evident by rewriting (1) as:

$$V(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + C(p, Q) \right\}$$

The unstructured models are statistical artifacts that allow the decision maker to assess formally the potential consequences of misspecification as captured by the construction of $C(\cdot, Q)$. In this paper we provide a formal interpretation of $C(\cdot, Q)$ as an index of misspecification fear: the lower the index, the higher the fear.[5]

We formalize model uncertainty by allowing the decision maker to posit a set Q. In our normative approach, it is then natural to enrich the standard decision-theoretic setting by taking Q as a given, a *datum* of the decision problem. For instance, in the climate policy problem, Q is the set of climate models that the policy maker considers. In this regard, observe that we are not after detecting which choice behavior of the decision maker may reveal model misspecification concerns, a different revealed preference exercise that would indeed require an endogenous Q.[6] In line with standard practice in applied economics, we imagine that the substantive modeling that underlies the construction of the elements of Q is simplified with an explicit structure imposed to facilitate interpretation. Applied researchers commonly avoid reducing model building to the construction of the complex black boxes that a purely nonparametric exercise might well involve, especially in multivariate settings.

[5] To ease terminology, we often refer to "misspecification" rather than "model misspecification."

[6] In this exercise, the findings of Denti and Pomatto (2022) may be useful.

A protective belt When c takes the entropic form $\lambda R(p||q)$, with $\lambda > 0$, criterion (1) takes the form

$$\min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + \lambda \min_{q\in Q} R(p||q)\right\} \tag{2}$$

proposed by Hansen and Sargent (2022). It is the most tractable version of criterion (1), which for a singleton Q further reduces to a standard multiplier criterion a la Hansen and Sargent (2001, 2008). By exchanging orders of minimization, we preserve this tractability and provide a revealing link to this earlier research,

$$\min_{q\in Q}\left\{\min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + \lambda R(p||q)\right\}\right\} \tag{3}$$

The inner minimization problem gives rise to the minimization problem featured by Hansen and Sargent (2001, 2008) to confront the potential misspecification of a given probability model q.[7] Unstructured models lack the substantive motivation of structured models, yet in (1) they act as a protective belt against model misspecification. The importance of their role is proportional – as quantified by λ in (2) – to their proximity to the set Q, a measure of their plausibility in view of the decision maker information. The outer minimization over structured models is the counterpart to the Wald (1950) and the more general Gilboa and Schmeidler (1989) max-min criterion. Clearly, the same inner-outer interpretation applies to the general case (1).

Our analysis provides a decision-theoretic underpinning for incorporating misspecification concerns in a distinct way from ambiguity aversion. Observe that misspecification fear is absent when the index $\min_{q\in Q} c(p,q)$ equals the indicator function δ_Q of the set of structured models Q, that is,

$$\min_{q\in Q} c(p,q) = \begin{cases} 0 & \text{if } p\in Q \\ +\infty & \text{else} \end{cases}$$

In this case, which corresponds to $\lambda = +\infty$ in (2), criterion (1) takes a max-min form:

$$V(f) = \min_{q\in Q}\int u(f)\,\mathrm{d}q$$

This max-min criterion thus characterizes decision makers who confront model misspecification but are not concerned by it, so are misspecification neutral (see Section 4.3). The criterion in (1) may thus be viewed as representing decision makers who use a more prudential variational criterion than if they were to max-minimize over the set of structured models which they posited. In particular, the farther away an unstructured model is from the set Q (so the less plausible it is), the less it is weighted in the minimization.

[7] The Hansen and Sargent (2001, 2008) formulation of preferences builds on extensive literature in control theory starting with Jacobson (1973)'s deterministic robustness criterion and a stochastic extension given by Petersen et al. (2000), among several others.

Axiomatics Our axiomatic analysis considers as in Gilboa et al. (2010) a mental preference $\succsim_Q^*$, describing the decision maker genuine ranking over acts, and a behavioral preference $\succsim_Q$ governing choices. The former is typically incomplete as the decision maker may find it difficult to rank all acts, the latter is instead complete because, at the end, a choice has to be made.

As it should become clear as our analysis unfolds, this modelling choice is conceptually important for the study of misspecification because it is the mental preference that, interestingly, turns out to be the one relevant for the analysis of misspecification attitudes. In particular, the flexibility of our two-preference setting allows us to capture misspecification sensitivity via a suitable weakening of the independence axiom, the weak certainty-independence axiom, for the mental preference. We show that this weak form of independence, which underlies variational representations, is needed to account for misspecification sensitivity. Indeed, stronger weakenings of the independence axiom, like the certainty-independence axiom, would force misspecification neutrality. These key decision-theoretic points, which underlie the importance of a two-preference setting, are discussed in detail in Sections 4.2 and 4.3.

Another key feature of our axiomatization is the use of a family $\mathcal{Q}$ of sets of models Q. We thus consider preferences $\succsim_Q^*$ and $\succsim_Q$ indexed by the sets Q and introduce axioms ensuring their consistency across different sets of models, each depicting a different possible information that the decision maker may have. In this way, our analysis is able to consistently connect different decision environments in which the decision maker may end up. Besides its inherent motivation, this rich setting also significantly eases the exposition. A derivation for a single given Q is, however, provided in the Online Appendix B.4.2. Appendix A uses the single-preference framework in axiomatizing the entropic criterion (2). We include this analysis to help situate our main criterion in a prior literature, including the single-preference variational model.

For an outline of our approach, let us consider the entropic case (2). Start with a singleton $Q = \{q\}$. Decision makers, being afraid that the reference model q might not be correct, contemplate also unstructured models $p \in \Delta$ and rank acts f according to the multiplier criterion

$$V_{\lambda,q}(f) = \min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + \lambda R(p||q)\right\} \tag{4}$$

Here the positive scalar λ is interpreted as an index of misspecification fear. When decision makers posit a nonsingleton set Q of structured models, but are concerned that none of them is correct, the multiplier criterion (4) then gives only an incomplete *dominance relation*:

$$f \succsim_Q^* g \iff V_{\lambda,q}(f) \geq V_{\lambda,q}(g) \qquad \forall q \in Q \tag{5}$$

With (5), decision makers can safely regard f better than g. Through this ranking, the dominance relation provides a preferential account of the probabilistic information that Q represents. The dominance relation thus naturally arises when the set Q is posited.

Yet, the ranking (5) has little traction because of the incomplete nature of $\succsim_Q^*$. Nonetheless, the burden of choice will have decision makers select between alternatives, be they rankable by $\succsim_Q^*$ or not. A cautious way to complete the binary relation $\succsim_Q^*$ is given by the preference $\succsim_Q$ represented by (2), or equivalently by (3). This criterion thus emerges in our analysis as a cautious completion of a multiplier dominance relation $\succsim_Q^*$. In this way, the probabilistic information gets embedded in the behavioral preference. Suitably extended to a general preference pair $(\succsim_Q^*, \succsim_Q)$, we support this approach by axiomatizing criterion (1) as the representation of the behavioral preference $\succsim_Q$ and the unanimity criterion

$$f \succsim_Q^* g \Longleftrightarrow \min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + c(p,q)\right\} \geq \min_{p\in\Delta}\left\{\int u(g)\,\mathrm{d}p + c(p,q)\right\} \qquad \forall q \in Q$$

as the representation of the incomplete dominance relation $\succsim_Q^*$.

To sum up, our two-preference approach is motivated by the natural way with which the dominance relation arises when the set Q is posited. In our approach, we connect the dominance and behavioral preferences to derive their desired representations. We then extend our analysis beyond the max-min case allowing for a more general non-variational criterion that encompasses a Bayesian formulation.

2 Preliminaries

2.1 Mathematics

Basic notions We consider a non-trivial algebra Σ of *events* in a *state space* S of payoff-relevant contingencies. We denote by Δ the set of finitely additive probabilities $p : \Sigma \to [0,1]$ and endow Δ and any of its subsets with the weak* topology (unless otherwise specified, these subsets are to be intended non-empty). Product sets are endowed with the product topology.

We denote by Δ^σ the subset of Δ formed by the countably additive probability measures. Given a probability q in Δ, we denote by $\Delta^\sigma(q)$ the collection of all probabilities p in Δ^σ that are absolutely continuous with respect to q, i.e., $q(A) = 0$ implies $p(A) = 0$ for all $A \in \Sigma$.

The (convex analysis) indicator function $\delta_C : \Delta \to [0,\infty]$ of a subset C of Δ is defined by

$$\delta_C(p) = \begin{cases} 0 & \text{if } p \in C \\ +\infty & \text{else} \end{cases}$$

Throughout we adopt the convention $0 \cdot \pm\infty = 0$.

Collections In what follows $\mathcal{Q}$ denotes a collection of compact subsets of Δ^σ. It is often assumed to be *proper*, that is, to contain all singletons and cover all doubletons.[8] Examples of

[8] That is, for each $q, q' \in \Delta^\sigma$ there exists some $Q \in \mathcal{Q}$ such that $\{q, q'\} \subseteq Q$.

proper collections $\mathcal{Q}$ are the collection of all finite sets of Δ^σ, the collection of all its compact subsets, the collection of all its polytopes as well as the collection $\mathcal{K}$ of its compact and convex subsets.

Statistical distances We say that a map $c : \Delta \times \Delta^\sigma \to [0, \infty]$ is a *statistical distance* if it is lower semicontinuous and satisfies the distance property

$$c(p, q) = 0 \Longleftrightarrow p = q$$

Given a statistical distance c and a family of compact sets $\mathcal{Q}$ in Δ^σ, we can define a *Hausdorff statistical set distance* $C : \Delta \times \mathcal{Q} \to [0, \infty]$ by

$$C(p, Q) = \min_{q \in Q} c(p, q)$$

It is easy to see that C is well defined, lower semicontinuous in the first argument and satisfies the following two properties:

(C.i) for each $Q \in \mathcal{Q}$,

$$C(p, Q) = 0 \Longleftrightarrow p \in Q$$

(C.ii) for each $Q, Q' \in \mathcal{Q}$,

$$Q \supseteq Q' \Longrightarrow C(\cdot, Q) \leq C(\cdot, Q')$$

These two properties make possible to interpret C as a set distance.

Divergences We say that a statistical distance $c : \Delta \times \Delta^\sigma \to [0, \infty]$ is a *divergence* if, for each $q \in \Delta^\sigma$,

$$c(p, q) < \infty \Longrightarrow p \ll q$$

Divergences thus assign an infinite penalty when p is not absolutely continuous with respect to q. To introduce a well-known class of divergences, assume that Σ is a σ-algebra. Given a continuous strictly convex function $\phi : [0, \infty) \to [0, \infty)$, with $\phi(1) = 0$ and $\lim_{t \to \infty} \phi(t)/t = +\infty$, define $D_\phi : \Delta \times \Delta^\sigma \to [0, \infty]$ by

$$D_\phi(p||q) = \begin{cases} \int \phi\left(\dfrac{\mathrm{d}p}{\mathrm{d}q}\right) \mathrm{d}q & \text{if } p \in \Delta^\sigma(q) \\ +\infty & \text{otherwise} \end{cases} \tag{6}$$

under the conventions $0/0 = 0$ and $\ln 0 = -\infty$. It can be proved that D_ϕ is a convex divergence, called *ϕ-divergence*.[9] When $\phi(t) = t \ln t - t + 1$, D_ϕ reduces to the relative entropy $R(p||q)$,

[9] For basic properties of ϕ-divergences we refer, for example, to Chapter 1 of Liese and Vajda (1987). As usual, $\mathrm{d}p/\mathrm{d}q$ denotes any version of the Radon-Nikodym derivative of p with respect to q.

while when $\phi(t) = (t-1)^2/2$ it becomes the Gini index $\chi^2(p||q)$.

Example Let Z be a metric space endowed with its Borel σ-algebra. Take $S = Z^\infty$ with the algebra Σ of cylinders. Given any $p \in \Delta$, we denote by p_t its restriction to the σ-algebra Σ_t of cylinders of length $t+1$. Define the *discounted relative entropy* $R_\beta : \Delta \times \Delta^\sigma \to [0,\infty]$ by

$$R_\beta(p||q) = (1-\beta)\sum_{t=0}^{\infty} \beta^t R(p_t||q_t)$$

where $\beta \in (0,1)$. Hansen and Sargent (2008) use this divergence when studying infinite-horizon discounted problems. It is routine to verify that R_β is a convex divergence. Notice that is possible for p_t to be absolutely continuous with respect to q_t for all t without p being absolutely continuous with respect to q over the σ-algebra generated by the algebra Σ of cylinders.[10]

2.2 Decision theory

Setup We consider a generalized Anscombe and Aumann (1963) setup where a decision maker chooses among uncertain alternatives described by (simple) acts $f : S \to X$, which are Σ-measurable simple functions from a *state space* S to a *consequence space* X. This latter space is assumed to be a non-empty convex set in a vector space like, for instance, the set of all (simple) lotteries defined on some prize space. The triple

$$(S, \Sigma, X) \tag{7}$$

forms an (Anscombe-Aumann) *decision framework*.

Let us denote by $\mathcal{F}$ the set of all acts. Given any consequence $x \in X$, we denote by $x \in \mathcal{F}$ also the constant act equal to x. Thus, under a standard abuse of notation we may identify X with the subset of constant acts in $\mathcal{F}$.

A preference $\succsim$ is a binary relation on $\mathcal{F}$ that satisfies the so-called *basic conditions* (cf. Gilboa et al., 2010), i.e., it is:

(i) *reflexive* and *transitive*;

(ii) *monotone*: for all $f, g \in \mathcal{F}$,

$$f(s) \succsim g(s) \quad \forall s \in S \Longrightarrow f \succsim g$$

[10] With an appropriate scaling, a limiting version as $\beta \uparrow 1$ converges to a divergence that is central to a discrete-time Donsker-Varadhan large deviation theory for ergodic Markov processes.

(iii) *c-continuous*: for all $x, y, z \in X$, the sets

$$\{\alpha \in [0,1] : \alpha x + (1-\alpha) y \succsim z\} \quad \text{and} \quad \{\alpha \in [0,1] : z \succsim \alpha x + (1-\alpha) y\}$$

are closed;

(iv) *non-trivial*: there exist $f, g \in \mathcal{F}$ such that $f \succ g$.

Moreover, a preference $\succsim$ is:

1. *continuous* if, for all $f, g, h \in \mathcal{F}$, the sets

$$\{\alpha \in [0,1] : \alpha f + (1-\alpha) g \succsim h\} \quad \text{and} \quad \{\alpha \in [0,1] : h \succsim \alpha f + (1-\alpha) g\}$$

are closed;

2. *unbounded* if, for each $x, y \in X$ with $x \succ y$, there exist $z, z' \in X$ such that

$$\frac{1}{2}z + \frac{1}{2}y \succsim x \succ y \succsim \frac{1}{2}x + \frac{1}{2}z'$$

Bets are binary acts that play a key role in decision theory. Formally, given any two prizes $x \succ y$, a bet on an event A is the act xAy defined by

$$xAy(s) = \begin{cases} x & \text{if } s \in A \\ y & \text{else} \end{cases}$$

In words, a bet on event A is a binary act that yields a more preferred consequence when A obtains.

Comparative uncertainty aversion Let $\succsim_1$ and $\succsim_2$ be two preferences on $\mathcal{F}$. As in Ghirardato and Marinacci (2002), we say that $\succsim_1$ is *more uncertainty averse than* $\succsim_2$ if, for each consequence $x \in X$ and act $f \in \mathcal{F}$,

$$f \succsim_1 x \implies f \succsim_2 x$$

In words, a preference is more uncertainty averse than another one if, whenever this preference is "bold enough" to prefer an uncertain alternative over a sure one, so does the other one.

An absolute notion of uncertainty aversion can be defined in this comparative setting once an uncertainty neutral preference is identified. In this case, a preference is declared to be uncertainty averse when more uncertainty averse than the neutral one.

Decision criteria We say that a complete preference $\succsim$ on $\mathcal{F}$ is *variational* when it is represented by a decision criterion $V : \mathcal{F} \to \mathbb{R}$ given by

$$V(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + c(p) \right\} \tag{8}$$

where the affine utility function $u : X \to \mathbb{R}$ is non-constant and the index of uncertainty aversion $c : \Delta \to [0, \infty]$ is grounded (i.e., $\min_{\Delta} c = 0$), lower semicontinuous and convex. In particular, given two unbounded variational preferences $\succsim_1$ and $\succsim_2$ on $\mathcal{F}$ that share the same u, we have that $\succsim_1$ is more uncertainty averse than $\succsim_2$ if and only if $c_1 \leq c_2$ (see Maccheroni et al., 2006, Propositions 6 and 8).

When the function c has the entropic form $c(p, q) = \lambda R(p||q)$ under a reference probability $q \in \Delta^{\sigma}$ and a coefficient $\lambda > 0$, criterion (8) takes the *multiplier* form

$$V_{\lambda, q}(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda R(p||q) \right\}$$

analyzed by Hansen and Sargent (2001, 2008).[11] If, instead, the function c has the indicator form δ_C, with C compact and convex, criterion (8) takes the *max-min* form

$$V(f) = \min_{p \in C} \int u(f) \, \mathrm{d}p$$

axiomatized by Gilboa and Schmeidler (1989).

All these criteria are here considered in their classical interpretation, so Waldean for the max-min criterion, in which the elements of Δ are interpreted as models.

3 Models and preferences

3.1 Models

The consequences of the acts among which decision makers have to choose depend on the exogenous states s that are outside their control. They know, however, that states obtain according to a probabilistic model described by a probability measure in Δ, the so-called *true* or *correct model*. If decision makers knew the true model, they would confront only risk, which is the randomness inherent to the probabilistic nature of the model. Our decision makers, unfortunately, may not know the true model. Yet, they are able to posit a set of *structured* probabilistic models Q, based on their information (possibly including existing scientific theories, say economic or physical), that forms a set of alternative hypotheses regarding the true model. It is

[11] Strzalecki (2011) provides the behavioral assumptions that characterize multiplier preferences among variational preferences.

a classical assumption, in the spirit of Wald (1950), in which Q is a set of posited hypotheses about the probabilistic behavior of a, natural or social, phenomenon of interest.

A *decision framework under model uncertainty* is described by a quartet:

$$(S, \Sigma, X, Q) \tag{9}$$

in which a set Q of models is added to a standard decision framework (7), as discussed in the Introduction. The true model might not be in Q, that is, the decision makers information may be unable to pin it down. Throughout the paper we assume that decision makers are aware of this limitation of their information and so confront model misspecification.[12] This is in contrast with Wald (1950) and most of the subsequent decision-theoretic literature, which assumes that decision makers either know the true model and so face risk or, at least, know that the true model belongs to Q and so face model ambiguity.[13]

Example (continued) We consider an example of a real investment problem with a single stochastic option for transferring goods from one period to another. This problem could be a planner's problem supporting a competitive equilibrium outcome associated with a stochastic growth model with a single capital good. We introduce an exogenous stochastic technology process that has an impact on the growth rate of capital as an example of what we call a structured model. This stochastic technology process captures what a previous literature in macro-finance has referred to as "long-run risk." For instance, see Bansal and Yaron (2004).[14]

We extend this formulation by introducing an unknown parameter θ used to index members of a parameterized family of stochastic technology processes. The investor's exploration of the entire family of these processes reflects uncertainty among possible structured models. We also allow the investor to entertain misspecification concerns over the parameterized models of the stochastic technology.[15] Later in this paper we will reconsider this example by allowing the investor to also be a statistician endowed with Bayesian priors over the parameter space Θ.

To represent this example formally, consider the Euclidean spaces W, Z and Θ, modelling respectively a random shock process, a stochastic technology (inclusive of a long-run risk component) and a parameter specification. We take the state space $S = Z^\infty$, endowed with the algebra $\Sigma = \bigcup_{t \geq 0} \Sigma_t$ of cylinders. A probability measure r on W^∞, known to the investor, captures risk in the economy. For instance, the implied process for the random shocks could be i.i.d. and distributed as a multivariate standard normal at each date t.

[12] Aydogan et al. (2023) propose an experimental setting that reveals the relevance of model misspecification for decision making.

[13] The model ambiguity (or uncertainty) literature is reviewed in Marinacci (2015).

[14] While this and some of the papers we cite, assume an endowment economy, their insights extend to a model with production.

[15] Hansen and Sargent (2021) give a (continuous time) example of such a specification with time varying parameters residing in a convex subset of an infinite-dimensional parameter space.

The exogenous (system) state vector z_t used to capture fluctuations in the technological opportunities has realizations in Z and the shock vector w_t has realizations in W. We build the exogenous technology process from the shocks in a parameter dependent way:

$$z_{t+1} = \psi\left(z_t, w_{t+1}, \theta\right) \tag{10}$$

for a given initial condition z_0. For instance, in long-run risk modeling one component of z_t evolves as a first-order autoregression:

$$z_{t+1}^1 = a_\theta z_t^1 + b_\theta^1 \cdot w_{t+1}$$

and another component is given by:

$$z_{t+1}^2 = d_\theta + b_\theta^2 \cdot w_{t+1}$$

At each time t the investor observes past and current values $\mathbf{z}^t = \{z_0, z_1, ..., z_t\}$ of the technology process, but does not know θ and does not directly observe the random shock vector w_t. The recursive formulation (10) implies, for a given observed $\mathbf{z}^t$, a mapping $\tau_\theta : W^\infty \to S = Z^\infty$ defined by

$$\tau_\theta\left((w_t)\right)_{t+1} = \begin{bmatrix} \sum_{s=0}^t a_\theta^{t-s} b_\theta^1 \cdot w_{s+1} + a_\theta^{t+1} z_0^1 \\ d_\theta + b_\theta^2 \cdot w_{t+1} \end{bmatrix} \qquad \forall t \geq 0$$

The probability r on W^∞ then induces for each θ a structured model q_θ on S via the pushforward distribution

$$q_\theta\left(C\right) = r\left(\tau_\theta^{-1}\left(C\right)\right)$$

for each cylinder C of S. As the shocks' distribution r is known, the parameter θ is the only source of model uncertainty. A nonsingleton parameter set Θ then translates in a nonsingleton set $Q = \{q_\theta\}_{\theta \in \Theta}$ of structured models.

Similarly, we consider a recursive representation of capital evolution given by

$$k_{t+1} = k_t \varphi\left(i_t/k_t, z_{t+1}\right)$$

where consumption $c_t \geq 0$ and investment $i_t \geq 0$ are constrained by an output relation:

$$c_t + i_t = \kappa k_t$$

for a pre-specified initial condition k_0. The parameter κ captures the productivity of capital. By design this technology is homogenous of degree one, which opens the door to stochastic growth as assumed in long-run risk models. Both i_t and c_t are constrained to be functions of $\mathbf{z}^t$ at each date t reflecting the observational constraint that θ is unknown to the investor in contrast to the history $\mathbf{z}^t$ of the technology process.[16] Formally, they are Σ_t-measurable functions.

Preferences are defined over consumption processes. Thus, the consequence space X consists of the simple lotteries defined over streams of consumption levels (c_t). An act f associates, to each realized sequence of the technology process, a lottery over streams of consumption levels restricted to depend on technology histories.[17] When such a lottery is degenerate, at each period t the act returns a consumption level c_t.

In this intertemporal setting, we consider an investor who solves a date 0 commitment problem. We pose this as a static problem with consumptions and investments that depend on information as it gets realized.[18] The affine utility function u over X is a discounted expected utility over lotteries. In particular, by considering degenerate lotteries, the utility of a consumption stream (c_t) is $(1-\beta)\sum_{t=0}^{\infty}\beta^t v(c_t)$, where $\beta \in (0,1)$ is a subjective discount rate and $v : \mathbb{R} \to \mathbb{R}$ is a utility function over consumption levels. The production technology further constraints the consumption and hence acts of interest, which form a collection C dependent on the initial condition k_0 for capital. These feasible acts f feature Σ_t-measurable sections f_t because of the observational constraints.

In a traditional analysis, the agent is assumed to know the true parameter θ^*, thus facing only risk. At the decision time $t = 0$, when only z_0 is known, the agent uses the standard expected utility objective function

$$\int_S u(f(s))\,\mathrm{d}q_{\theta^*}(s) = (1-\beta)\sum_{t=0}^{\infty}\beta^t \int_S v(f_t(s))\,\mathrm{d}q_{\theta^*}(s) \tag{11}$$

to solve the decision problem

$$\max_f \int u(f)\,\mathrm{d}q_{\theta^*} \quad \text{sub}\, f \in C \tag{12}$$

Here the agent confronts risk via a singleton set $Q = \{q_{\theta^*}\}$ consisting of the true model, used in an expected utility criterion. Yet, typically agents do not know the true model and confront model uncertainty via a nonsingleton set Q of structured models. In the rest of the section we present a preferential analysis of a rational agent coping with model uncertainty, yielding the decision criterion discussed in the Introduction that extends traditional expected utility analysis under risk to model uncertainty.

[16] Note that the endogenous state variable, k_t, reveals no new information in addition to current and past values of the technology process. This means that there is no incentive for the investor to experiment in this setting.

[17] This shows that the elements of Z^∞ are the payoff-relevant contingencies, which motivates our choice of $S = Z^\infty$ as the state space.

[18] By posing this as a date 0 commitment problem, we deliberately avoid dynamic consistency considerations.

3.2 Preferences

We consider a two-preference setup as in Gilboa et al. (2010). In our first axiomatization, we allow the set of structured models to vary within a collection of sets of models $\mathcal{Q}$ which we assume to be proper (as defined in Section 2.1). In the second approach, developed in the Online Appendix, Q will be fixed.

For each $Q \in \mathcal{Q}$ we thus consider a mental preference $\succsim_Q^*$ and a behavioral preference $\succsim_Q$.

Definition 1 *A preference $\succsim_Q$ is (*subjectively*)* rational *if it is:*

a. complete;

b. solvable: for each $f \in \mathcal{F}$ there exists $x \in X$ such that $f \sim_Q x$;[19]

c. risk independent: for all $x, y, z \in X$ and $\alpha \in (0,1)$,

$$x \sim_Q y \Longrightarrow \alpha x + (1-\alpha)\, z \sim_Q \alpha y + (1-\alpha)\, z$$

The behavioral preference $\succsim_Q$ governs the decision maker choice behavior and so it is natural to require it to be complete as, eventually, the decision maker has to choose between alternatives (burden of choice). It is subjectively rational because, in an "argumentative" perspective, the decision maker cannot be convinced that it leads to incorrect choices. Risk independence ensures that $\succsim_Q$ is represented on the space of consequences X by an affine utility function $u : X \to \mathbb{R}$, for instance an expected utility functional when X is the set of lotteries. So, risk is addressed in a standard way and we abstract from non-expected utility issues.

The mental preference $\succsim_Q^*$ on $\mathcal{F}$ represents the decision maker "genuine" preference over acts, so it has the nature of a dominance relation for the decision maker. As such, it might well not be complete because of the decision maker inability to compare some pairs of acts. These preferences have an antecedent in statistical decision theory in the study of admissibility, where decision rules are evaluated using a partial ordering based on their *ex ante* performance over a family of possible parameter values.

Definition 2 *A continuous preference $\succsim_Q^*$ is a* dominance relation *(or is* objectively rational*) if it is:*

a. c-complete: for all $x, y \in X$, $x \succsim_Q^* y$ or $y \succsim_Q^* x$;

b. completeness: $\succsim_Q^*$ is complete when Q is a singleton;

c. weak c-independent: for all $f, g \in \mathcal{F}$, $x, y \in X$ and $\alpha \in (0,1)$,

$$\alpha f + (1-\alpha)x \succsim_Q^* \alpha g + (1-\alpha)x \implies \alpha f + (1-\alpha)y \succsim_Q^* \alpha g + (1-\alpha)y$$

[19] As well-known, a complete continuous preference relation that satisfies risk independence is solvable.

d. convex: for all $f, g, h \in \mathcal{F}$ and $\alpha \in (0,1)$,

$$f \succsim_Q^* h \text{ and } g \succsim_Q^* h \implies \alpha f + (1-\alpha) g \succsim_Q^* h$$

When $f \succsim_Q^* g$ we say that f *dominates* g. The dominance relation is objectively rational because the decision maker can convince others of its reasonableness, for instance through arguments based on scientific theories (a case especially relevant for our purposes). Momentarily, axiom A.3 will further clarify the nature of the dominance relation. Axiomatically, it is a variational preference not required to be complete, unless Q is a singleton.[20] In the singleton Q case the dominance relation is complete and yet, because of model misspecification, satisfies only a weak form of independence. Hence in our approach, model misspecification may cause violations of the independence axiom for the dominance relation. Later in the paper, Proposition 4 will show that relaxing independence to weak c-independence is conceptually necessary as, otherwise, the behavioral preference would be misspecification neutral. This is a key decision-theoretic observation for our analysis.

By adding preferences $\succsim_Q^*$ and $\succsim_Q$ to (9) we form a *two-preference decision environment under model uncertainty*

$$\left(S, \Sigma, X, Q, \succsim_Q^*, \succsim_Q\right) \tag{13}$$

The next two assumptions, which we take from Gilboa et al. (2010), connect the two preferences $\succsim_Q^*$ and $\succsim_Q$.

A.1 *Consistency*: for all $f, g \in \mathcal{F}$,

$$f \succsim_Q^* g \Longrightarrow f \succsim_Q g$$

Consistency asserts that, whenever possible, the mental ranking informs the behavioral one. The next condition, caution, says that the decision maker opts by default for a sure alternative x over an uncertain one f, unless the dominance relation says otherwise.

A.2 *Caution*: for all $x \in X$ and $f \in \mathcal{F}$,

$$f \not\succsim_Q^* x \Longrightarrow x \succsim_Q f$$

Unlike the previous assumptions, the next two are peculiar to our analysis. They both link the posited set Q to the two preferences $\succsim_Q^*$ and $\succsim_Q$ of the decision maker. We begin with the

[20] Convexity is stronger than uncertainty aversion a la Schmeidler (1989), which merely requires that $f \sim_Q^* g$ implies $\alpha f + (1-\alpha) g \succsim_Q^* g$ for all $\alpha \in (0,1)$. Yet, convexity and uncertainty aversion coincide under completeness (see, e.g., Lemma 56 of Cerreia-Vioglio et al., 2011). Nascimento and Riella (2011) study incomplete variational preferences, but their result is not applicable to our setting because their axioms are over lotteries of acts (and their state space is finite).

dominance relation $\succsim_Q^*$. Here we write $f \overset{Q}{=} g$ when $q(f = g) = 1$ for all $q \in Q$, i.e., f and g are equal almost everywhere according to each structured model.

A.3 *Objective Q-coherence*: for all $f, g \in \mathcal{F}$,

$$f \overset{Q}{=} g \implies f \sim_Q^* g$$

This axiom provides a preferential translation of the special status of structured models over unstructured ones: if they all regard two acts to be almost surely identical, the decision maker's "genuine" preference $\succsim_Q^*$ follows suit and ranks them indifferent.

In what follows we will see that, though the set Q of structured models might not be convex per se, its closed convex hull $\overline{\text{co}}Q$ that contains "hybrid models" might be of interest.[21] This is also mirrored in our next axiom. Even when Q is not convex, we assign a special role to the probabilities in its convex hull relative to other unstructured models. Our rationale is that hybrid models retain an epistemic status and are more than just statistical artifacts used to assess model misspecification.

To introduce our next axiom, recall that a rational preference $\succsim_Q$ satisfies risk independence and thus admits an affine utility function $u : X \to \mathbb{R}$ over consequences.[22] Given a model $p \in \Delta$ and an act f, we define an indifference class $X_f^p \subseteq X$ of consequences x_f^p via the equality

$$u(x_f^p) = \int u(f)\, \mathrm{d}p \tag{14}$$

We can interpret each x_f^p as a consequence that would be indifferent, so equivalent, to act f if p were the correct model. By constructing these equivalent consequences for alternative acts and models, our next axiom relates the posited set of models Q with the behavioral preference $\succsim_Q$.

A.4 *Subjective Q-coherence*: for all $f \in \mathcal{F}$ and $x \in X$,

$$x \succ_Q^* x_f^p \implies x \succ_Q f$$

if and only if $p \in \overline{\text{co}}Q$.

To motivate the right-hand side of the implication imposed in this axiom, by construction the decision maker would be indifferent between the act f and a constant counterpart x_f^p under expected utility with probability p. This construction is achieved without any misspecification

[21] Loosely speaking, hybrid models are probabilities obtained as mixture of structured models. Subsequently, we will suggest a robust Bayesian perspective that can justify this convexity.

[22] We make the usual identification of constant acts with consequences. Though in principle u might depend on Q, in our analysis it will turn out to be constant across Q's.

concerns regarding the probability p. On the other hand, misspecification concerns potentially reduce the attractiveness of f relative to x_f^p when p is in the convex hull of the set of structured probability models Q. This axiom presumes that this implication only applies to such p's. This salience of p is the preferential footprint of a structured or hybrid model that the decision maker takes seriously under consideration because of its epistemic status – as opposed to a purely unstructured model, which they regard as a mere statistical artifact with no epistemic content.

While we state Axioms A.3 and A.4 using a set of structured probabilities Q, as we will see, Corollary 3 shows that our main results hold even when these axioms are required to hold only for singletons $\{q\}$. When Q is fixed and does not range over a class $\mathcal{Q}$, the general versions of these two coherence axioms are needed for the analysis. We use this fixed Q approach for an analysis of the entropic formulation of model misspecification aversion in Appendix A and for the general formulation considered in Online Appendix B.4.2. For this reason, we stated the two coherence axioms for a generic Q.

To conclude, observe that in traditional purely subjective axiomatizations there is no way – actually, no language – to embed the probabilistic information that Q represents in the decision maker preference.[23] The last two axioms provide the needed embedding, as the representation theorems will show momentarily.

Example (continued) At the end of the one-sector growth example discussed above, the agent confronted risk via a singleton set $Q = \{q_{\theta^*}\}$ consisting of the true model, used in the expected utility criterion (11). When the agent still entertains a single model, but now has doubts about it being correct, our approach prescribes that $\succsim_Q^*$ is a complete variational preference and $\succsim_Q$ coincides with it. In this case, we still have a singleton $Q = \{q_\theta\}$ but we dropped the star since the investor no longer knows whether the single structured model is correct. As a consequence, the preference $\succsim_Q^*$ considers other probabilities besides q_θ, but they are penalized by a cost function. The Q-coherence axioms discipline such a penalization as we will discuss below. For example, $\succsim_Q^*$ could be described by the discounted entropic criterion

$$V_{\theta,\beta}(f) = \min_{p \in \Delta^\sigma} \left\{ \int u(f)\, \mathrm{d}p + \lambda R_\beta(p||q_\theta) \right\}$$

When $Q = \{q_\theta\}_{\theta\in\Theta}$ is not a singleton, $\succsim_Q^*$ is not complete and has, as it will be seen in both our two-preference axiomatizations, a multi-variational form. For example,

$$f \succsim_Q^* g \iff V_{\theta,\beta}(f) \geq V_{\theta,\beta}(g) \qquad \forall \theta \in \Theta \tag{15}$$

[23] For instance, in the Gilboa and Schmeidler (1989) seminal axiomatization the derived set of probabilities C is purely subjective. There is no formal connection with any underlying probabilistic information, something left to the decision maker personal, unmodelled, elaborations. A notable exception is Gajdos et al. (2008), which considers probabilistic information. Its analysis proceeds along lines very different from ours.

Consistency and caution help to connect the evaluations that represent $\succsim_Q^*$ to the decision criterion that represents $\succsim_Q$. Specifically, consistency imposes that when each evaluation of $\succsim_Q^*$ deems f better than g, so does $\succsim_Q$. As we will see, it implies that $\succsim_Q$ is represented by a rule that aggregates these evaluations (cf. Proposition 7). We will also show that caution singles out the "min" as the aggregation rule (cf. Theorem 1). Thus, when $Q = \{q_\theta\}_{\theta\in\Theta}$ is compact, the agent decision problem (12) becomes

$$\max_f \min_{p\in\Delta^\sigma} \left\{ \int u(f)\,\mathrm{d}p + \lambda \min_{\theta\in\Theta} R_\beta(p||q_\theta) \right\} \quad \text{sub}\, f \in C$$

The dominance relation (15) proves to be useful in solving this problem by ruling out all strongly dominated acts.[24]

Parameters and utility functions Many applications in econometrics and statistics use loss functions (negatives of utility functions) expressed directly in terms of unknown parameters that can be inferred in part from data. In our view, this is often done to short circuit the process of specifying a substantively interesting application of a decision problem in which, like in our example, the unknown parameters have implications for objects of interest such as consumption or policy outcomes. That said, with some modest reframing, unknown parameters could be included as arguments in the utility function in an extended analysis.

4 Decision criteria and model misspecification

4.1 Main criterion

We introduced a two-preference decision environment under model uncertainty (13) as a tuple

$$\left(S, \Sigma, X, Q, \succsim_Q^*, \succsim_Q\right)$$

with the dependence of preferences on Q highlighted. Decision environments, however, may share common state and consequence spaces, but differ on the posited sets of structured models because of the different information that decision makers may have. It then becomes important to ensure that decision makers use decision criteria that, across such environments, are consistent.

To provide our first foundation of our main decision criterion, in this section we consider a family

$$\left\{\left(S, \Sigma, X, Q, \succsim_Q^*, \succsim_Q\right)\right\}_{Q\in\mathcal{Q}}$$

of decision environments that differ in the set Q of posited models, which vary in a collection

[24] That is, all acts $g \in C$ for which there exist $f \in C$ and $\varepsilon > 0$ with $V_{\theta,\beta}(f) \geq V_{\theta,\beta}(g) + \varepsilon$ for all $\theta \in \Theta$.

$\mathcal{Q}$ that we continue to assume to be proper. Next we introduce three axioms on the family $\{\succsim_Q^*\}_{Q\in\mathcal{Q}}$ that connect these environments.

A.5 *Monotonicity* (*in model ambiguity*): for all $f,g\in\mathcal{F}$ and all $Q'\subseteq Q$,

$$f \succsim_Q^* g \Longrightarrow f \succsim_{Q'}^* g$$

According to this axiom, when the "structured" information underlying a set Q is good enough for the decision maker to establish that an act dominates another one, a better information which decreases model ambiguity can only confirm such judgement. Its reversal would be, indeed, at odds with the objective rationality spirit of the dominance relation.

Next we consider a separability assumption.

A.6 *Q-separability*: for all $f,g\in\mathcal{F}$,

$$f \succsim_q^* g \quad \forall q\in Q \implies f \succsim_Q^* g$$

In words, an act dominates another one when it does, separately, through the lenses of each structured model. In this axiom the incompleteness of $\succsim_Q^*$ arises as that of a Paretian order over the, complete but possibly misspecification averse, preferences $\succsim_q^*$ determined by the elements of Q.

We close with a continuity axiom. To state it, we need a last piece of notation: we denote by $x_{f,q}$ the consequence indifferent to act f for the preference $\succsim_q^*$.[25]

A.7 *Lower semicontinuity*: for all $x\in X$ and $f\in\mathcal{F}$, the set $\{q\in\Delta^\sigma : x \succsim_q^* x_{f,q}\}$ is closed.

Next we introduce a class

$$P_{\mathcal{Q}} = \{(\succsim_Q^*, \succsim_Q)\}_{Q\in\mathcal{Q}}$$

of two-preference families that builds on the properties that we have introduced.

Definition 3 *A two-preference family $P_{\mathcal{Q}}$ is (*misspecification*) *robust *if:*

(i) $\{\succsim_Q^*\}_{Q\in\mathcal{Q}}$ is monotone, separable and lower semicontinuous;

(ii) for each $Q\in\mathcal{Q}$, $\succsim_Q^*$ is an unbounded dominance relation, $\succsim_Q$ is a rational preference, both are Q-coherent and jointly satisfy caution and consistency.

We can now state our first representation result.

Theorem 1 *Let $P_{\mathcal{Q}}$ be a two-preference family. The following statements are equivalent:*

[25] In symbols, $f \sim_q^* x_{f,q}$. In particular, $x_{f,q}$ should not be confused with x_f^q as in (14).

(i) $P_{\mathcal{Q}}$ *is robust;*

(ii) there exist an onto affine function $u : X \to \mathbb{R}$ *and a divergence* $c : \Delta \times \Delta^{\sigma} \to [0, \infty]$, *convex in* p, *such that, for each* $Q \in \mathcal{Q}$,

$$f \succsim_Q^* g \Longleftrightarrow \min_{p\in\Delta} \left\{ \int u(f)\,\mathrm{d}p + c(p,q) \right\} \geq \min_{p\in\Delta} \left\{ \int u(g)\,\mathrm{d}p + c(p,q) \right\} \qquad \forall q \in Q \tag{16}$$

and

$$f \succsim_Q g \Longleftrightarrow \min_{p\in\Delta} \left\{ \int u(f)\,\mathrm{d}p + \min_{q\in Q} c(p,q) \right\} \geq \min_{p\in\Delta} \left\{ \int u(g)\,\mathrm{d}p + \min_{q\in Q} c(p,q) \right\} \tag{17}$$

for all acts $f, g \in \mathcal{F}$.

Moreover, u *is cardinal and, given* u, c *is unique.*

A robust $P_{\mathcal{Q}}$ is thus characterized by a utility and divergence pair (u, c) that, consistently across decision environments, represents each $\succsim_Q^*$ via the unanimity rule (16) and each $\succsim_Q$ via the decision criterion

$$V_Q(f) = \min_{p\in\Delta} \left\{ \int u(f)\,\mathrm{d}p + \min_{q\in Q} c(p,q) \right\} \tag{18}$$

The Hausdorff statistical set distance

$$p \mapsto C(p, Q) = \min_{q\in Q} c(p,q)$$

between p and Q is strictly positive if and only if p is an unstructured model, i.e., $p \notin Q$. In particular, the more distant from Q is an unstructured model, the more it is penalized and so the smaller is its role in the minimization problem that criterion (18) features. An unstructured model p may play a role in this criterion when $c(p,q) < \infty$ for some structured model q, that is, when it has a finite distance from a structured model. Momentarily, we will engage in a comparative analysis of misspecification aversion that will permit us to interpret C as an index of misspecification aversion.

A few remarks are now in order. Before moving to them, observe that in this representation theorem there is no convexity assumption on the sets of structured models. In Section 4.4, we will study the convex case.

Unstructured penalization Objective Q-coherence has important implications for the cost functions used to impose misspecification aversion. Consider first the case where Q is a singleton $\{q\}$. Objective Q-coherence imposes that all the unstructured models which are not absolutely continuous with respect to q are infinitely penalized. In contrast, subjective Q-coherence prescribes that all the unstructured models are penalized, i.e., $c(p,q) > 0$ whenever $p \neq q$.

When Q is not a singleton, the absolute continuity restriction entailed by objective Q-coherence continues to apply when all the probabilities in Q are mutually absolutely continuous. Though it is a restriction that we do not make, it is commonly imposed in statistical problems when constructing a likelihood function. For such problems we may think of the potential misspecification of each model as a way to represent a misspecified likelihood function.[26] Finally, subjective Q-coherence imposes that all the unstructured models are strictly penalized.

Admissibility Since the dominance relation $\succsim_Q^*$ requires unanimity across structured models as stated in (16), it implies a counterpart to admissibility extended to accommodate misspecification aversion. Analogous to the outcome from the standard formulation of statistical decision theory, the partial ordering $\succsim_Q^*$ alone suffices to rule out a collection of acts as potential solutions to decision problems.[27]

Specifications and computability Two specifications of our representation are noteworthy. First, when c is the entropic statistical distance $\lambda R(p||q)$, with $\lambda \in (0, \infty]$, we have the following important tractable version of our representation

$$V_Q(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda \min_{q \in Q} R(p||q) \right\} \tag{19}$$

Specifically, for $\lambda \in (0, \infty)$,[28]

$$\min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda \min_{q \in Q} R(p||q) \right\} = \min_{q \in Q} -\lambda \log \int e^{-\frac{u(f)}{\lambda}} \mathrm{d}q \tag{20}$$

This result is well known when Q is a singleton, that is, when (19) is a standard multiplier criterion.

A second noteworthy special case of our representation is the Gini criterion

$$V_Q(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda \min_{q \in Q} \chi^2(p||q) \right\} \tag{21}$$

Remarkably, we have

$$\min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda \min_{q \in Q} \chi^2(p||q) \right\} = \min_{q \in Q} \left\{ \int u(f) \, \mathrm{d}q - \frac{1}{2\lambda} \mathrm{Var}_q(u(f)) \right\} \tag{22}$$

for all acts f for which the *mean-variance* (in utils) criteria on the r.h.s. are monotone. So,

[26] Removing objective Q-coherence from Theorem 1 is equivalent to requiring c to be solely a statistical distance and not necessarily a divergence. A similar observation applies to Proposition 7.

[27] We further elaborate in the working paper version.

[28] When $\lambda = \infty$, we have $\min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda \min_{q \in Q} R(p||q) \right\} = \min_{q \in Q} \int u(f) \, \mathrm{d}q$. See Online Appendix B.4.1 for the simple proof of (20).

the Gini criterion is a monotone version of the max-min mean-variance criterion.[29]

As to computability, in the important case when criterion (18) features a ϕ-divergence, like the specifications just discussed, we need only to know the set Q to compute it, no integral with respect to unstructured models is needed. This is proved in the next result, a consequence of a duality formula of Ben-Tal and Teboulle (2007).[30]

Proposition 1 *Let $Q \subseteq \Delta^\sigma$ and $\lambda \in (0, \infty)$. If Q is compact, for each act $f \in \mathcal{F}$,*

$$V_Q(f) = \min_{p \in \Delta} \left\{ \int u(f) \, \mathrm{d}p + \lambda \min_{q \in Q} D_\phi(p||q) \right\} = \lambda \min_{q \in Q} \sup_{\eta \in \mathbb{R}} \left\{ \eta - \int \phi^* \left(\eta - \frac{u(f)}{\lambda} \right) \mathrm{d}q \right\}$$

By using integrals only under structured models, the r.h.s. formula substantially simplifies computations and thus confirms the analytical tractability of the previous specifications.

4.2 Comparative misspecification aversion

We now turn to the study of comparative misspecification aversion. From an axiomatic point of view, we identified nonneutrality toward model misspecification with the dominance relation $\succsim_Q^*$ violating the independence axiom and only satisfying weak c-independence. This suggests that, in our approach, model misspecification aversion is captured by the dominance relation $\succsim_Q^*$.

Definition 4 *A robust two-preference family $P_{1,\mathcal{Q}}$ is* more (model) misspecification averse *at Q than $P_{2,\mathcal{Q}}$ if, for each $x \in X$ and $f \in \mathcal{F}$,*

$$f \succsim_{1,Q}^* x \Longrightarrow f \succsim_{2,Q}^* x$$

We say that $P_{1,\mathcal{Q}}$ is more misspecification averse *than $P_{2,\mathcal{Q}}$ if this implication holds for all $Q \in \mathcal{Q}$.*

In words, a decision maker is more misspecification averse at Q when her dominance relation is more uncertainty averse. In a similar way we define *more uncertainty aversion (at Q)* by replacing $\succsim_{i,Q}^*$ with $\succsim_{i,Q}$ for $i = 1, 2$.

Throughout the section, the two decision makers' preferences are represented by the pairs (u_1, c_1) and (u_2, c_2) identified in Theorem 1. With this, we can now state an important equivalence result.

Proposition 2 *Let $P_{1,\mathcal{Q}}$ and $P_{2,\mathcal{Q}}$ be robust two-preference families and $Q \in \mathcal{Q}$. The following statements are equivalent:*

[29] At the end of Online Appendix B.4.1 we further discuss this point.

[30] Here ϕ^* denotes the convex Fenchel conjugate of ϕ, once extended to $\mathbb{R}$ by setting $\phi(t) = +\infty$ if $t < 0$. In particular, ϕ^* is real valued and increasing.

(i) $P_{1,\mathcal{Q}}$ is more misspecification averse at Q than $P_{2,\mathcal{Q}}$;

(ii) $P_{1,\mathcal{Q}}$ is more uncertainty averse at Q than $P_{2,\mathcal{Q}}$.

When the maps $p \mapsto C_1(p,Q)$ and $p \mapsto C_2(p,Q)$ are convex, this is equivalent to

(iii) u_1 is cardinally equivalent to u_2 and $C_1(\cdot,Q) \leq C_2(\cdot,Q)$, provided $u_1 = u_2$.

This characterization immediately yields that higher aversion to misspecification is equivalent to higher uncertainty aversion. Functionally it translates, under a minor convexity assumption,[31] into a lower statistical set distance. For this reason, in our main representation we may interpret

$$C(\cdot,Q) = \min_{q\in Q} c(\cdot,q)$$

as an index of misspecification aversion at each Q.

In our main criterion uncertainty attitudes are thus equated to misspecification attitudes, something that will not happen for the less extreme criteria that will be discussed in Section 5. To understand why this is the case, we further elaborate on the equivalence between points (i) and (ii) above. The consistency assumption yields

$$f \succsim_Q^* x \implies f \succsim_Q x \tag{23}$$

that can be read as saying that the mental preference is more uncertainty averse than the behavioral preference one. At the same time, by the continuity of $\succsim_Q^*$ and $\succsim_Q$, the caution assumption can be rewritten as

$$f \succsim_Q x \implies f \succsim_Q^* x \tag{24}$$

that we can read in the opposite way. As a consequence, in our main model $\succsim_Q^*$ and $\succsim_Q$ share the same uncertainty aversion attitudes. Since in our main representation the preference $\succsim_Q$ turns out to be variational, this immediately yields the inequality $C_1(\cdot,Q) \leq C_2(\cdot,Q)$.

This observation allows us to better understand how misspecification aversion affects uncertainty aversion. Let us consider again relation (23). Intuitively, the extra uncertainty aversion of $\succsim_Q^*$ can be ascribed to two factors: a genuine extra aversion of $\succsim_Q^*$ or its incompleteness. Indeed, consider a consequence x and an act f with $f \succsim_Q x$ and $f \not\succsim_Q^* x$. The fact that x is not preferred to f by $\succsim_Q^*$ can happen because either $x \succ_Q^* f$ or x and f are not comparable. To further elaborate, assume that either Q is finite or Q is the convex hull of a finite set and c is convex. Consistency then rules out the first possibility, i.e., $x \succ_Q^* f$. In fact, even for our general model (cf. Proposition 7), $\succsim_Q^*$ admits a Paretian representation as in (16). Given our assumptions on Q and the properties of c, this would imply the existence of a consequence y

[31] Used only to prove that (ii) implies (iii).

such that $x \succ_Q^* y \succ_Q^* f$. In turn, consistency would yield that $x \succ_Q y \succsim_Q f$, which is not compatible with $f \succsim_Q x$. Therefore, even without caution, the extra uncertainty aversion featured by $\succsim_Q^*$ is due to its incompleteness, which in turn follows from model ambiguity (Q not being a singleton). Thus, the lower uncertainty aversion featured by $\succsim_Q$ relative to $\succsim_Q^*$ can be imputed to how this incompleteness is resolved, in other words to the decision makers attitudes toward model ambiguity. Functionally, under caution for each act f the worst evaluation given by $\succsim_Q^*$ is the one followed by $\succsim_Q$. In such an extreme case, $\succsim_Q$ is as uncertainty averse as $\succsim_Q^*$, as we have already seen preferentially. Later in the paper, Section 5 will discuss less extreme criteria.

In our last result, we characterize comparative uncertainty attitudes at a global level.

Corollary 1 *Let $P_{1,\mathcal{Q}}$ and $P_{2,\mathcal{Q}}$ be robust two-preference families. The following statements are equivalent:*

(i) $P_{1,\mathcal{Q}}$ is more misspecification averse than $P_{2,\mathcal{Q}}$;

(ii) $P_{1,\mathcal{Q}}$ is more uncertainty averse than $P_{2,\mathcal{Q}}$;

(iii) u_1 is cardinally equivalent to u_2 and $c_1 \leq c_2$, provided $u_1 = u_2$;

(iv) u_1 is cardinally equivalent to u_2 and $C_1 \leq C_2$, provided $u_1 = u_2$.

4.3 Misspecification neutrality

As it should be clear by now, it is the dominance relation $\succsim_Q^*$ that captures misspecification attitudes. It is then natural to expect that misspecification neutrality should be a notion that pertains to $\succsim_Q^*$. At the same time, we just learned that in our main criterion the uncertainty attitudes of $\succsim_Q^*$ and $\succsim_Q$ coincide, so one might want to discuss misspecification neutrality also at the level of $\succsim_Q$. We thus have three different approaches: (i) an axiomatic one for $\succsim_Q^*$, (ii) a functional one for $\succsim_Q$ (discussed in Online Appendix), (iii) a "combo" one for $\succsim_Q$. Next we discuss each of them and show that, remarkably, they lead to the same conclusions. Besides its own interest, this can be seen as a consistency check for our analysis.

4.3.1 Axiomatic approach for $\succsim_Q^*$

In our analysis, we identified the presence of model misspecification concerns with violations of the independence axiom by $\succsim_Q^*$. This prompts us to the following definition.

Definition 5 *Let $P_\mathcal{Q}$ be a robust two-preference family and $Q \in \mathcal{Q}$. The preference $\succsim_Q^*$ is (*model*) *misspecification neutral* at Q if it satisfies independence.*

We next show that misspecification neutrality leads to the models in Q being fully trusted by both $\succsim_Q^*$ and $\succsim_Q$.

Proposition 3 *Let $P_{\mathcal{Q}}$ be a robust two-preference family and $Q \in \mathcal{Q}$. The following statements are equivalent:*

(i) $\succsim^*_Q$ *is misspecification neutral at* Q*;*

(ii) for each $f, g \in \mathcal{F}$,

$$f \succsim^*_Q g \iff \int u(f)\,\mathrm{d}q \geq \int u(g)\,\mathrm{d}q \qquad \forall q \in Q \tag{25}$$

In this case, we have, for each $f, g \in \mathcal{F}$,

$$f \succsim_Q g \iff \min_{q \in Q} \int u(f)\,\mathrm{d}q \geq \min_{q \in Q} \int u(g)\,\mathrm{d}q \tag{26}$$

This result shows that when $\succsim^*_Q$ satisfies independence, the models are fully trusted and, in turn, the behavioral preference becomes Waldean. In other words, the uncertainty aversion featured by $\succsim_Q$ is just the result of model ambiguity and misspecification neutrality leads to the max-min Waldean criterion.

Not to have all models fully trusted, we therefore need to weaken independence. The next result will show that moving to weak c-independence is a necessary step. To discuss this key point, we need to introduce a classical weakening of independence which is stronger than weak c-independence.

A.8 *C-independence*: for all $f \in \mathcal{F}$, $x, y \in X$ and all $\alpha \in (0, 1]$,

$$f \succsim^*_Q x \iff \alpha f + (1-\alpha)\,y \succsim^*_Q \alpha x + (1-\alpha)\,y$$

Clearly, this axiom can also be stated for the preference $\succsim_Q$. When $\succsim_Q$ is a rational preference or the dominance relation $\succsim^*_Q$ is complete, our version of this axiom is equivalent to the original one of Gilboa and Schmeidler (1989). Otherwise, ours is weaker.

Proposition 4 *Let $P_{\mathcal{Q}}$ be a robust two-preference family and $Q \in \mathcal{Q}$. The following statements are equivalent:*

(i) $\succsim^*_Q$ *satisfies c-independence;*

(ii) $\succsim_Q$ *satisfies c-independence;*

(iii) for each $f, g \in \mathcal{F}$,

$$f \succsim_Q g \iff \min_{q \in Q} \int u(f)\,\mathrm{d}q \geq \min_{q \in Q} \int u(g)\,\mathrm{d}q$$

(iv) $\delta_{\overline{\text{co}}Q} \leq C(\cdot, Q) \leq \delta_Q$*. In particular,* $C(\cdot, Q) = \delta_Q$ *when* Q *is convex;*

(v) $c(p, q) = \infty$ *for all* $p \notin \overline{\text{co}}Q$ *and all* $q \in Q$*.*

Weakening independence to c-independence would thus lead to a behavioral preference that still fully trusts the models in Q, as point (iii) shows. From a statistical distance angle, this suggests that misspecification neutrality is the attitude of a decision maker who confronts model misspecification, but does not care about it: all the unstructured models that are not hybrid are infinitely penalized, as points (iv) and (v) indicate.

This angle becomes relevant here because it also shows that the representation (16) of the dominance relation becomes

$$f \succsim_Q^* g \Longleftrightarrow \min_{q' \in \overline{\text{co}}Q} \left\{ \int u(f)\, \mathrm{d}q' + c(q', q) \right\} \geq \min_{q' \in \overline{\text{co}}Q} \left\{ \int u(g)\, \mathrm{d}q' + c(q', q) \right\} \qquad \forall q \in Q$$

Unstructured models play no role here. Only structured and hybrid models are relevant.

4.3.2 Combo approach for $\succsim_Q$

As misspecification aversion arises when structured models are not trusted, the following notion that combines functional and preferential ingredients seems natural.

Definition 6 *Let* P_Q *be a robust two-preference family and* $Q \in \mathcal{Q}$*. The preference* $\succsim_Q$ *is (*model*)* misspecification neutral *at* Q *if*

$$\int u(f)\, \mathrm{d}q \geq \int u(g)\, \mathrm{d}q \quad \forall q \in Q \Longrightarrow f \succsim_Q g$$

for all $f, g \in \mathcal{F}$*.*

Here the decision maker trusts models enough so to follow them when they unanimously rank pairs of acts. Fear of misspecification thus becomes decision-theoretically irrelevant. For this reason, we classify this decision maker as model misspecification neutral. The next result shows that this neutral attitude characterizes a decision maker who adopts the max-min criterion (26).

Theorem 2 *Let* P_Q *be a robust two-preference family and* $Q \in \mathcal{Q}$*. The preference* $\succsim_Q$ *is misspecification neutral at* Q *if and only if it is represented by the max-min criterion (26).*

It is easy to see that the misspecification neutrality of $\succsim_Q^*$ at Q implies that of $\succsim_Q$. At a global level they become equivalent, as the next result shows.

Corollary 2 *Let* P_Q *be a robust two-preference family. The following statements are equivalent:*

(i) $\succsim_Q^*$ *is misspecification neutral at all* $Q \in \mathcal{Q}$*;*

(ii) $\succsim_Q$ *is misspecification neutral at all* $Q \in \mathcal{Q}$*;*

(iii) $c(p,q) = \delta_{\{q\}}(p)$ *for all* $p \in \Delta$ *and for all* $q \in \Delta^\sigma$*.*

In this case, we have, for each $Q \in \mathcal{Q}$*,*

$$f \succsim_Q^* g \iff \int u(f)\,\mathrm{d}q \geq \int u(g)\,\mathrm{d}q \quad \forall q \in Q \tag{27}$$

and

$$f \succsim_Q g \iff \min_{q \in Q} \int u(f)\,\mathrm{d}q \geq \min_{q \in Q} \int u(g)\,\mathrm{d}q \tag{28}$$

4.3.3 Discussion

These results provide the sought-after decision-theoretic argument for the interpretation of the max-min criterion as the special case of our decision criterion (18) that corresponds to aversion to model ambiguity, with no fear of misspecification. As remarked in the Introduction, under this interpretation a decision maker using criterion (18) may be viewed as a decision maker who, under model ambiguity, would max-minimize over the set of structured models which she posited but that, for fear of misspecification, ends up using the more prudential variational criterion (18). Unstructured models lack the informational status of structured models, yet in criterion (18) they act as a "protective belt" against model misspecification.

In particular, the special multiplier case of a singleton $Q = \{q\}$ then corresponds to a decision maker who, with no fear of misspecification, would adopt the expected utility criterion $\int u(f)\,\mathrm{d}q$ to confront the risk inherent to q. In other words, a singleton Q in (18) corresponds to an expected utility decision maker who fears misspecification.

Summing up, in our analysis decision makers adopt the max-min criterion (28) when they either confront only model ambiguity (an information trait) or are averse to model ambiguity with no fear of model misspecification (a taste trait).

4.3.4 Misspecification aversion (absolute)

Having identified misspecification neutrality at Q of $\succsim_Q^*$ and $\succsim_Q$ respectively with the expected utility dominance relation (27) and the Waldean criterion (28), we may declare a robust two-preference family $P_{\mathcal{Q}}$ *misspecification averse at* Q when either $\succsim_Q^*$ is more uncertainty averse than the dominance relation (27) or $\succsim_Q$ is more uncertainty averse than the Waldean criterion (28). No matter which choice we make, a robust two-preference family $P_{\mathcal{Q}}$ is misspecification averse.

That said, we conclude this subsection by studying a mild form of misspecification aversion: models are trusted in some specific cases. To this end, note that structured models may be incorrect, yet useful as Box (1976) famously remarked. This motivates the next notion. Recall that act xAy, with $x \succ_Q y$, represents a bet on event A.

Definition 7 *A preference* $\succsim_Q$ *is* bet-consistent *if, given any* $x \succ_Q y$,

$$q(A) \geq q(B) \quad \forall q \in Q \Longrightarrow xAy \succsim_Q xBy$$

for all events $A, B \in \Sigma$.

Under bet-consistency, a decision maker may fear model misspecification, yet regards structured models as good enough to choose to bet on events that they unanimously rank as more likely. Preferences that are bet-consistent can be classified as exhibiting a mild form of fear of model misspecification. The following result shows that an important class of preferences, ones for which the cost specification is a scaled ϕ-divergence, are bet consistent.

Proposition 5 *If* $\lambda \in (0, \infty]$ *and* $c = \lambda D_\phi$, *then a preference* $\succsim_Q$ *represented by (18) is bet-consistent.*

This result applies to criterion (19) as a special case. More generally, it sheds light on the decision-theoretic nature of the tractable specifications of our criterion based on ϕ-divergences.

4.4 Convex sets of models

In this final subsection we sharpen Theorem 1 by assuming that the sets of models are compact and convex. To do so, we first need to discuss the role of convexity.

We previously encountered a closed convex hull of a set of models in the statement of axiom A.4 and in the discussion that followed. Conceptually it is not an innocuous operation: a hybrid model that mixes two structured models can only be less well motivated than either of them. Decision criterion (18) accounts for the lower appeal of hybrid models when c is convex, like for instance when it is a ϕ-divergence. To see why, observe that $\min_{p \in \Delta} \left\{ \int u(f)\, \mathrm{d}p + c(p, q) \right\}$ is, for each act f, convex in q. In turn, this implies that hybrid models negatively affect criterion (18). This negative impact of mixing thus features an "aversion to model hybridization" attitude, behaviorally captured by axiom A.9 below. Remarkably, the relative entropy criterion turns out to be neutral to model hybridization. In this important special case, convexity of Q plays a little role (as Online Appendix B.4.1 clarifies).

The convexification of Q can be also justified by building a convex family of probability distributions from a set of "structured building block" or primitive models weighted by possible prior distributions. The convexity can then be imposed on the set of priors used in the weighting. For instance, each of the primitive models could each have an i.i.d. representation. By entertaining a prior weighting over these we obtain an exchangeable process. As is known from the Hewitt and Savage (1955) version of the de Finetti Representation Theorem, conversely we may represent any exchangeable process as probability weighted average of i.i.d. processes. By entertaining uncertainty about the weighting captured by a convex set of prior distributions, we can in this way obtain a convex specification of Q. Incorporating misspecification concerns provides a protective shield for each of the resulting exchangeable processes. In Section 6, we will describe alternative extensions of our analysis that allow for conceptually distinct ways to confront model misspecification and prior uncertainty.

We introduce a new axiom based on this added convexity structure on sets of models (it features the same terminology of axiom A.7). Observe that under the hypotheses of Theorem 1, all dominance relations $\succsim_Q^*$ agree on X and so we can just write $\succsim^*$, dropping the subscript Q.

A.9 *Model hybridization aversion*: for all $q, q' \in \Delta^\sigma$, $\lambda \in (0,1)$ and $f \in \mathcal{F}$,

$$\lambda x_{f,q} + (1-\lambda)\, x_{f,q'} \succsim^* x_{f,\lambda q + (1-\lambda) q'}$$

According to this axiom, the decision maker dislikes, *ceteris paribus*, facing a hybrid structured model $\lambda q + (1-\lambda)\, q'$ that, by mixing two structured models q and q', could only have a less substantive motivation.

The next result extends Theorem 1 by dealing with sets of structured models that are also convex; in particular, here we get a convex divergence. Recall that $\mathcal{K}$ is the proper collection of compact and convex sets.

Proposition 6 *Let $P_{\mathcal{K}}$ be a two-preference family. The following statements are equivalent:*

(i) $P_{\mathcal{K}}$ is robust and model hybridization averse;

(ii) there exist an onto affine $u : X \to \mathbb{R}$ and a convex divergence $c : \Delta \times \Delta^\sigma \to [0, \infty]$ such that, for each $Q \in \mathcal{K}$,

$$f \succsim_Q^* g \iff \min_{p \in \Delta} \left\{ \int u(f)\, \mathrm{d}p + c(p,q) \right\} \geq \min_{p \in \Delta} \left\{ \int u(g)\, \mathrm{d}p + c(p,q) \right\} \qquad \forall q \in Q$$

and

$$f \succsim_Q g \iff \min_{p \in \Delta} \left\{ \int u(f)\, \mathrm{d}p + \min_{q \in Q} c(p,q) \right\} \geq \min_{p \in \Delta} \left\{ \int u(g)\, \mathrm{d}p + \min_{q \in Q} c(p,q) \right\}$$

for all acts $f, g \in \mathcal{F}$.

Moreover, u is cardinal and, given u, c is unique.

5 Beyond caution

As caution is the axiom behind the prudential nature of our representation result, it is natural to wonder about what happens when we dispense with it. To this end we introduce a new class of two-preference families.

Definition 8 *A two-preference family $P_{\mathcal{Q}}$ is (*misspecification*)* sensitive *if:*

(i) $\left\{\succsim_Q^*\right\}_{Q\in\mathcal{Q}}$ *is monotone, separable and lower semicontinuous;*

(ii) for each $Q \in \mathcal{Q}$, $\succsim_Q^*$ is an unbounded dominance relation, $\succsim_Q$ is a rational preference, both are Q-coherent when restricted to singletons and jointly satisfy consistency.

Compared to the notion of robust family (Definition 3), we made two changes. The important one is the removal of caution. We also require Q-coherence to hold only when Q is a singleton, a change immaterial under caution as we will later discuss (we could have actually considered this weaker version throughout). As a result, given a sensitive two-preference family $P_{\mathcal{Q}}$ and a set of models Q, the dominance relation continues to be represented as follows:

$$f \succsim_Q^* g \iff \min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + c(p,q)\right\} \geq \min_{p\in\Delta}\left\{\int u(g)\,\mathrm{d}p + c(p,q)\right\} \quad \forall q \in Q$$

In particular, an act f induces an evaluation map

$$q \mapsto \min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + c(p,q)\right\} \tag{29}$$

over the collection Q of structured models. Our criterion (18) emerges when these evaluations are aggregated via the minimum on Q. But, in principle, less extreme stances are conceivable. This requires dropping caution, as the next result shows.[32]

Proposition 7 *Let $P_{\mathcal{Q}}$ be a two-preference family. The following statements are equivalent:*

(i) $P_{\mathcal{Q}}$ is sensitive;

(ii) there exist an onto affine $u : X \to \mathbb{R}$, a divergence $c : \Delta \times \Delta^\sigma \to [0,\infty]$, convex in p, and for each $Q \in \mathcal{Q}$ a normalized and monotone functional $J_Q : B(Q) \to \mathbb{R}$ such that

$$f \succsim_Q^* g \iff \min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + c(p,q)\right\} \geq \min_{p\in\Delta}\left\{\int u(g)\,\mathrm{d}p + c(p,q)\right\} \qquad \forall q \in Q \tag{30}$$

and

$$f \succsim_Q g \iff J_Q\left(\min_{p\in\Delta}\left\{\int u(f)\,\mathrm{d}p + c(p,\cdot)\right\}\right) \geq J_Q\left(\min_{p\in\Delta}\left\{\int u(g)\,\mathrm{d}p + c(p,\cdot)\right\}\right) \tag{31}$$

for all acts $f, g \in \mathcal{F}$.

Moreover, u is cardinal and, given u, c is unique.

[32] $B(Q)$ is the space of all real-valued bounded Borel measurable functions with domain Q.

Decision-theoretically, Theorem 1 is the special case of this result when $\succsim_Q^*$ and $\succsim_Q$ jointly satisfy caution for all $Q \in \mathcal{Q}$, as Corollary 3 will momentarily show. Analytically, it corresponds to the special case where J_Q is the minimum over Q of the evaluation maps (29). In this case, by exchanging the order of minima, (31) reduces to the decision criterion (18). We will explore an altogether different case of J_Q in the next section, for instance a quasi-arithmetic specification.

Corollary 3 *Let $P_\mathcal{Q}$ be a two-preference family. The following statements are equivalent:*

(i) $P_\mathcal{Q}$ is robust;

(ii) $P_\mathcal{Q}$ is sensitive and $(\succsim_Q^*, \succsim_Q)$ jointly satisfy caution for all $Q \in \mathcal{Q}$;

(iii) there exist an onto affine $u : X \to \mathbb{R}$ and a divergence $c : \Delta \times \Delta^\sigma \to [0, \infty]$, convex in p, such that (16) and (17) hold for all $Q \in \mathcal{Q}$.

Moreover, u is cardinal and, given u, c is unique.

Given the equivalence between points (i) and (ii), this corollary shows that for our main results we only need to consider the Q-coherence Axioms A.3 and A.4 when restricted to singletons, $\{q\} \in \mathcal{Q}$. Notice in particular that $p \in \overline{\text{co}}\,\{q\}$ in Axiom A.4 simplifies to $p = q$, thus making q the only model that satisfies the implication that Axiom A.4 features.

6 A Bayesian analysis

With the exception of the Bayesian construction of convex sets of models (Section 4.4), our analysis so far has been conducted in a classical Waldean setting where the decision maker specifies a family of structured models of interest. In contrast, in this section we outline a Bayesian analysis based on prior uncertainty.

Under model ambiguity, the decision maker has, possibly multiple, prior probabilities μ_Q over the set of structured models Q. Typically, each such prior probability $\mu_Q(q)$ of a structured model $q \in Q$ quantifies the decision maker belief that q is the correct model. Under model misspecification, this interpretation is no longer possible because the decision maker no longer regards the correct probability model to be among the structured models. Thus, they no longer form an exhaustive collection of mutually exclusive uncertain alternatives. Nevertheless, the family of structured models continues to play a central role in the decision theory leaving the door open to imposing subjective priors, μ_Q, over these models.

In this section, we consider two approaches. One approach follows the Bayesian approach with a single prior, but entertains preferences for which the uncertainty induced by the prior is distinct from that contributed by risk. A second entertains multiple priors as in robust Bayesian analysis. In both cases, potential model misspecification continues to play a central role in our analysis. Interestingly, in his conclusion, Chamberlain (2020) emphasized the importance of sensitivity over both likelihoods and priors. He was led to a very special case of what follows as he chose to remain within previous decision theory under uncertainty, as we will comment below.

6.1 A smooth Bayesian criterion

We first consider a functional J_Q in criterion (31) that is a quasi-arithmetic mean over the evaluation maps (29):

$$V_Q(f) = \phi_Q^{-1}\left(\int_Q \phi_Q\left(\min_{p\in\Delta}\left\{\int_S u(f(s))\,\mathrm{d}p(s) + c(p,q)\right\}\right) d\mu_Q(q)\right) \tag{32}$$

This is, formally, a Bayesian criterion with the prior probability μ_Q interpreted as an averaging device over the structured models. In this representation, the variational criteria indexed by Q are

$$\min_{p\in\Delta}\left\{\int_S u(f)\,\mathrm{d}p + c(p,q)\right\}$$

They account for fear of misspecification about the posited models q, while the function ϕ_Q addresses the fear of prior misspecification. The Bayesian criterion (32) incorporates model misspecification concerns into the smooth ambiguity criterion of Klibanoff et al. (2005), which is the special case $c(p,q) = \delta_{\{q\}}(p)$ that imposes model misspecification neutrality. In this regard, observe that our comparative uncertainty aversion analysis extends to this more general setting (in particular, the equivalence in Proposition 3).

An important entropic specification of criterion (32) is

$$V_Q^{\lambda,\xi}(f) = \phi_\xi^{-1}\left(\int_Q \phi_\xi\left(\min_{p\in\Delta}\left\{\int_S u(f(s))\,\mathrm{d}p(s) + \lambda R(p||q)\right\}\right) d\mu_Q(q)\right) \tag{33}$$

where $\phi_\xi(t) = -e^{-\frac{1}{\xi}t}$. The parameter $\xi > 0$ captures aversion to prior uncertainty, while the parameter $\lambda > 0$ is a fear of model misspecification index. The lower λ is the more misspecification averse is $\succsim_Q^*$ in the sense of Definition 4. Next we show that, as fear of either model or prior misspecification vanishes or explodes, we get the criteria that one would expect. This provides an analytical consistency check for criterion (33). Recently, Lanzani (2023) used criterion (35) to study learning under model misspecification.

Proposition 8 *Let* $\operatorname{supp}\mu_Q = Q$*. For each* $f \in \mathcal{F}$*,*

$$\lim_{\xi\to 0^+} V_Q^{\lambda,\xi}(f) = \min_{p\in\Delta}\left\{\int_S u(f)\,\mathrm{d}p + \lambda\min_{q\in Q} R(p||q)\right\} \quad \forall\lambda > 0 \tag{34}$$

and

$$\lim_{\xi\to\infty} V_Q^{\lambda,\xi}(f) = \int_Q \left(\min_{p\in\Delta} \left\{ \int_S u(f(s))\,\mathrm{d}p(s) + \lambda R(p||q) \right\} \right) d\mu_Q(q) \quad \forall \lambda > 0 \tag{35}$$

Moreover,

$$\lim_{\xi\to\infty}\lim_{\lambda\to\infty} V_Q^{\lambda,\xi}(f) = \lim_{\lambda\to\infty}\lim_{\xi\to\infty} V_Q^{\lambda,\xi}(f) = \int_Q \left(\int_S u(f(s))\,\mathrm{d}q(s) \right) d\mu_Q(q) \tag{36}$$

In words, the limit (34) shows that, as fear of prior misspecification explodes, criterion (33) gets closer and closer to our criterion (19). In contrast, the limit (35) shows that, when such fear vanishes, we end up with a criterion that averages, via the prior μ_Q, multiplier criteria (one per structured model q). Finally, the limit (36) shows that, when both fears vanish, at the limit we have the two-stage subjective expected utility criterion.[33] In deriving this result, we focused on the entropic formulation. But the result can be generalized in different directions, for example, by replacing either the relative entropy with a general divergence as in (6) or the conditions on ξ with similar ones on the Arrow-Pratt index of ϕ_Q.

6.2 A variational robust Bayesian criterion

Criterion (33) turns out to have an alternative interpretation as the reduced form of a preference criterion that incorporates a robust prior concern. Specifically, a generalization of criterion (33) is the outcome of the minimization over ν in:

$$V_Q(f) = \min_{\nu\ll\mu} \left\{ \int_Q \min_{p\in\Delta} \left\{ \int_S u(f(s))\,\mathrm{d}p(s) + c(p,q) \right\} d\nu(q) + \xi R(\nu||\mu_Q) \right\}$$

This has a direct extension to a specification with two divergences, one that captures the aversion to model misspecification and another that depicts aversion to prior misspecification, as proposed by Hansen and Sargent (2023) but without complete axiomatic support. Consistent with Proposition 7, one may replace $\xi R(\nu||\mu_Q)$ with a generic penalty function $d(\nu)$, leading to a robust Bayesian criterion

$$V_Q(f) = \min_{\nu} \left\{ \int_Q \min_{p\in\Delta} \left\{ \int_S u(f(s))\,\mathrm{d}p(s) + c(p,q) \right\} d\nu(q) + d(\nu) \right\} \tag{37}$$

This variational criterion is robust toward prior misspecification because of the minimization over ν. For instance, when d is the (convex analysis) indicator function of a compact set Γ of priors, criterion (37) takes the multiple-prior form à la Gilboa-Schmeidler

$$V_Q(f) = \min_{\nu\in\Gamma} \int_Q \min_{p\in\Delta} \left\{ \int_S u(f(s))\,\mathrm{d}p(s) + c(p,q) \right\} d\nu(q)$$

[33] As discussed in Cerreia-Vioglio et al. (2013).

Compared to a traditional robust Bayesian approach, criterion (37) takes into account also model misspecification via the inner term $\min_{p \in \Delta} \left\{ \int_S u(f(s))\, \mathrm{d}p(s) + c(p,q) \right\}$, which replaces the traditional term $\int_S u(f(s))\, \mathrm{d}q(s)$. Criterion (32) specializes to criterion (33) when a relative entropy penalty specification is used for both aversion to prior misspecification and to model misspecification. When the two relative entropy penalty parameters are equal, we have the preferences suggested by Chamberlain (2020) in his concluding section. That said, we leave a full-fledged analysis of these criteria and of their relationships to future research. We close by observing that if we take Γ to be the set of all possible priors, then the last multiple-prior criterion essentially collapses to the Waldean criterion studied earlier in the paper. This allows for interpreting the Waldean criterion as capturing a maximal specification of aversion to prior misspecification.

6.3 Example (concluded)

We now reconsider and complete our running example by incorporating a robust Bayesian perspective in which the parameter vector θ becomes learnable over time with observations on the stochastically evolving technology process given a prior distribution over Θ. When making decisions at date t, the investor can use observations on past and current values $\mathbf{z}^t = \{z_0, z_1, ..., z_t\}$ of the technology process to make inferences about θ. Other signal observable to the investor could also be included in the computations. For a given prior, there is typically a separation between prediction and control, meaning that recursive solution to Bayesian learning can be employed while adjusting the objective to accommodate potential misspecification for each of the possible models.

Collin-Dufresne et al. (2016) and Andrei et al. (2019) have explored learning implications in models with long-run risk and a unique prior. As was argued in these papers and in an earlier Hansen (2007) contribution, some of the technology parameters may be hard to estimate making the prior an important input in the calculations. When the posterior distributions are highly sensitivity to priors for extensive period of time, prior uncertainty becomes an important consideration for an investor.

By exploiting some well known feature of min-max optimization, robust prior adjustments can be implemented with computational approaches that iterate between minimized priors given a decision process and maximizing decisions computed with a given prior. A full exploration of such computational methods is beyond the scope of this paper, although they have been used in other settings. The point here is that our axiomatic formulations open the door to exploring misspecification of models and priors, where the latter would seem particular important for applications where the data is not sufficient to narrow substantially the scope of prior uncertainty. The extensive literature on partial identification provides another setting where, even asymptotically, prior uncertainty remains a concern in decision making.

6.4 On the interpretation of priors

As we previously remarked, under model misspecification a set Q of structured models is no longer a set of exhaustive and mutually exclusive alternatives, so a logical partition upon which to define a prior probability. What might be a new partition of this kind?

To address this question, denote by $p^* \in \Delta$ the correct model. The decision makers do not know whether or not it belongs to Q. Let q^* be the structured model, assumed to uniquely exist, such that

$$c(p^*, q^*) = \min_{q \in Q} c(p^*, q)$$

Model q^* best approximates, or best fits, the correct model p^* according to the statistical distance c that decision makers adopt. Under model ambiguity, when they know that p^* is in Q, we have $p^* = q^*$ and so q^* itself is the correct model.

Decision makers are uncertain about q^*, that is, about which structured model $q \in Q$ best fits the correct model. But, they know that one of them is, indeed, the best fit. Under this interpretation of its elements, Q thus forms a collection of exhaustive and mutually exclusive alternatives. Decision makers now regard each element q of Q as a "candidate best fitting model": this is how they interpret q and what they are uncertain about. The meaning of prior $\mu_Q(q)$ is then clear: it quantifies the decision maker belief that q is the best fit of the correct model (see Walker, 2013, for an insightful discussion).

This interpretation of μ_Q reduces to the standard one under model ambiguity because, as previously remarked, in this case the best fit coincides with the correct model itself. In the working paper version, we make more rigorous this discussion.

7 Conclusion

Quantitative researchers use models to enhance their understanding of economic phenomena and to make policy assessments. In essence, each model tells its own quantitative story. We refer to such models as "structured models." Typically, there are more than just one such type of model, with each giving rise to a different quantitative story. Statistical and economic decision theories have addressed how best to confront the ambiguity among structured models. Such structured models are, by their very nature, misspecified. Nevertheless, the decision maker seeks to use such models in sensible ways. This problem is well recognized by applied researchers, but it is typically not part of formal decision theory. In this paper, we extend decision theory to confront model misspecification concerns. In so doing, we recover a variational representation of preferences that includes penalization based on discrepancy measures between "unstructured alternatives" and the set of structured probability models.

A natural generalization of our main criterion is

$$V(f) = \min_{p \in \Delta} \left\{ \int u(f)\,\mathrm{d}p + C(p, Q) \right\}$$

where C is a general statistical set distance, not necessarily Hausdorff (so not necessarily characterized by an underlying statistical distance). This variational criterion still leads to a preference which is uncertainty averse. Though the analysis of this general criterion is beyond the scope of this paper and left for future research, we close our exposition with it, as its form should help to put our exercise in a final perspective. A further topic left for future research is a proper axiomatic analysis of the robust Bayesian criteria that we discussed in the last section.

A Entropic misspecification: a single preference approach

As discussed at length in the paper, we regard our two-preference approach as conceptually appropriate for the study of decision making under model misspecification. Yet, as mentioned in the Introduction, to help situate our main criterion (1) in the broader literature and to better relate it with the single-preference variational model, in this appendix we now develop a single-preference derivation of an entropic version (19) of our main criterion. This also shows that our modelling choice was not dictated by technical impediments (we focus on the entropic special case because of the exemplary nature of this appendix).

Specifically, we axiomatize the representation

$$V_{\lambda,Q}(f) = \min_{p \in \Delta} \left\{ \int u(f)\,\mathrm{d}p + \lambda \min_{q \in Q} R(p||q) \right\} \tag{38}$$

by using only the subjectively rational preference $\succsim$ on $\mathcal{F}$. The starting point is the following rewriting of (38), contained in Online Appendix B.4.1,

$$V_{\lambda,Q}(f) = \min_{q \in Q} \phi_\lambda^{-1} \left(\int \phi_\lambda(u(f))\,\mathrm{d}q \right)$$

with $\phi_\lambda(t) = -e^{-\frac{1}{\lambda}t}$ for all $t \in \mathbb{R}$ and $\lambda > 0$. Thus, to axiomatize (38) one can split the proof in a few parts:

1. provide an axiomatization for the criterion

$$V(f) = \min_{q' \in Q'} \phi^{-1} \left(\int \phi(u(f))\,\mathrm{d}q' \right) \tag{39}$$

where $Q' \subseteq \Delta^\sigma$ is compact and has an essential event,[34] $u : X \to \mathbb{R}$ is non-constant and

[34] An event $E \in \Sigma$ is *essential* when $\min_{q' \in Q'} q'(E) \in (0, 1)$.

affine, and $\phi : \operatorname{Im} u \to \mathbb{R}$ is a strictly increasing and continuous function;

2. prove that ϕ is concave;

3. prove that ϕ is CARA.

This strategy mirrors the one used by Strzalecki (2011) to axiomatize multiplier preferences, which correspond to a singleton Q'. An extra part is, however, needed in our setting:

4. show that $Q' = Q$ under a single-preference version of subjective Q-coherence.

Point 1 is achieved in Lemma 3. Other axiomatizations are available in the literature for criterion (39), which decision theoretically is a subjective version of the max-min criterion and mathematically a minimum of quasi-arithmetic means (see, e.g., Casadesus-Masanell et al., 2000 or Alon and Schmeidler, 2014). Points 2 and 3 are achieved in Lemma 1 by showing that ϕ is concave (resp. CARA) if and only if $\succsim$ is convex (resp. satisfies weak c-independence). Finally, point 4 is established in Lemma 2 where we show that the endogenous/subjective set Q' of point 1 coincides with the posited Q under the following single-preference version of subjective Q-coherence.

A.4* *Single-preference Subjective Q-coherence*: for all $f \in \mathcal{F}$ and $x \in X$,

$$x \succ x_f^p \Longrightarrow x \succ f$$

if and only if $p \in \overline{\text{co}}Q$.

Compared to the original version of the axiom we require also the first ranking, $x \succ x_f^p$, to be in terms of $\succsim$ as the dominance relation is here missing.

Lemma 1 *Let $\succsim$ be a preference represented by V defined as in (39) with an essential event. The following facts are true:*

1. *$\succsim$ satisfies convexity if and only if ϕ is concave;*

2. *$\succsim$ satisfies weak c-independence if and only if ϕ is CARA.*

The next lemma makes formal the "coincidence" of Q and Q' under single-preference subjective Q-coherence.

Lemma 2 *Let $\succsim$ be a preference represented by V defined as in (39) with an essential event. If $\succsim$ satisfies weak c-independence and convexity, then $\succsim$ satisfies single-preference subjective Q-coherence if and only if $\overline{\text{co}}Q = \overline{\text{co}}Q'$. In particular, Q' in (39) can be replaced with Q.*

To continue the analysis, we need a richer setting than the one used for our main results. In particular, we require Σ to be a σ-algebra, not just an algebra, and we assume that X is a convex subset of a topological vector space, not just a vector space. A preference $\succsim$ is *biseparable* when it is represented by a subcontinuous[35] functional $V : \mathcal{F} \to \mathbb{R}$ that features a non-constant restriction v on X and a capacity $\rho : \Sigma \to [0,1]$, with $\rho(E) \in (0,1)$ for some $E \in \Sigma$, such that

$$V(xAy) = v(x)\rho(A) + v(y)(1 - \rho(A))$$

for all $A \in \Sigma$ and for all $x, y \in X$ with $v(x) \geq v(y)$. Here an event E of Σ is essential when $x \succ xEy \succ y$ for all (some) $x \succ y$ in X; in this case $\rho(E) \in (0,1)$. In the next result $c_{xEy} \in X$ is the certainty equivalent of xEy, i.e., $c_{xEy} \sim xEy$.

Proposition 9 (Ghirardato et al., 2003) *Let $\succsim$ be a biseparable preference with representation V, $x, y, z \in X$ and E an essential event. The following statements are equivalent:*

(i) $x \succsim z \succsim y$ *and* $xEy \sim c_{xEz}Ec_{zEy}$*;*

(ii) $v(x) \geq v(y)$ *and* $v(z) = \dfrac{v(x) + v(y)}{2}$.

We call z a *preference midpoint* of x and y. It depends on x, y and $\succsim$, but neither on the representation V nor on the essential event E. If $x \succsim y$, such a z always exists because $\operatorname{Im} v$ is an interval; following Hardy, Littlewood, and Polya (1934), we denote any midpoint of x and y by $\mathfrak{M}(x,y)$. When $y \succsim x$, we still use $\mathfrak{M}(x,y)$ to denote any preference midpoint of y and x.

Preference midpoints of a pair (x,y) are typically not unique, but they form an equivalence class under $\sim$. Given any $f, g \in \mathcal{F}$, we denote by $\mathfrak{M}(f,g)$ any act $h \in \mathcal{F}$ such that $h(s) \sim \mathfrak{M}(f(s), g(s))$ for all $s \in S$. With this, Ghirardato et al. (2003) introduce the following axioms:

A.10 *Invariance*: for all $f, g \in \mathcal{F}$ and all $x \in X$, $f \succsim g$ if and only if $\mathfrak{M}(f,x) \succsim \mathfrak{M}(g,x)$.

A.11 *Ambiguity aversion*: for all $f, g \in \mathcal{F}$, $f \succsim g$ implies $\mathfrak{M}(f,g) \succsim g$.

A.12 *Ambiguity neutrality*: for all $f, g \in \mathcal{F}$, $f \succsim g$ implies $f \succsim \mathfrak{M}(f,g) \succsim g$.

A final standard axiom is needed to guarantee that probabilities are countably additive.

A.13 *Monotone continuity*: for all $x, y, z \in X$ with $y \succ z$ and all sequences of events $\{E_n\}_{n \geq 1} \subseteq \Sigma$ with $E_n \downarrow \emptyset$, there exists $\bar{n} \in \mathbb{N}$ such that $xE_{\bar{n}}y \succ z$.

We can now state the representation result achieving point 1.

[35] That is, $V(f_\eta) \to V(f)$ where f_η is a net in $\mathcal{F}$ that pointwise converges to $f \in \mathcal{F}$, with each f_η measurable with respect to the same finite Σ-measurable partition of S.

Lemma 3 *Let $\succsim$ be a binary relation on $\mathcal{F}$. The following statements are equivalent:*

(i) *$\succsim$ is a monotone continuous rational preference that satisfies biseparability, invariance, and ambiguity aversion.*

(ii) *there exist a non-constant, continuous and affine $u : X \to \mathbb{R}$, a strictly increasing and continuous function $\phi : \operatorname{Im} u \to \mathbb{R}$, and a compact set Q' in Δ^σ such that*

$$V(f) = \min_{q' \in Q'} \phi^{-1}\left(\int \phi(u(f))\,\mathrm{d}q'\right)$$

represents $\succsim$ on $\mathcal{F}$, with $0 < \min_{q' \in Q'} q'(E) < 1$ for some $E \in \Sigma$.

By merging Lemma 3 with Lemmas 1 and 2, we get the sought-after representation result. Observe that a continuous rational preference that satisfies weak c-independence and convexity is, axiomatically, a variational preference; if it also satisfies c-independence, it is a Gilboa-Schmeidler max-min preference.

Theorem 3 *Let $\succsim$ be a binary relation on $\mathcal{F}$. The following statements are equivalent:*

(i) *$\succsim$ is a monotone continuous rational preference that satisfies weak c-independence, convexity, biseparability, invariance, ambiguity aversion, and single-preference subjective Q-coherence;*

(ii) *there exist a non-constant, continuous and affine $u : X \to \mathbb{R}$ and $\lambda \in (0, \infty]$ such that*

$$V(f) = \min_{p \in \Delta}\left\{\int u(f)\,\mathrm{d}p + \lambda \min_{q \in Q} R(p||q)\right\}$$

represents $\succsim$ on $\mathcal{F}$.

Moreover,

1. *u is unique up to a positive affine transformation and, given u, λ is unique.*
2. *$\succsim$ is a multiplier preference, i.e., $Q' = \{q'\}$, if and only if it satisfies ambiguity neutrality.*
3. *$\succsim$ is a max-min preference, i.e., $\lambda = \infty$, if and only if it satisfies c-independence.*
4. *$\succsim$ is an expected utility preference, i.e., $Q' = \{q'\}$ and $\lambda = \infty$, if and only if it satisfies ambiguity neutrality and c-independence.*

References

[1] S. Alon and D. Schmeidler, Purely subjective maxmin expected utility, *Journal of Economic Theory*, 152, 382–412, 2014.

[2] D. Andrei, M. Hasler and A. Jeanneret, Asset pricing with persistence risk, *Review of Financial Studies,* 32, 2809–2849, 2019.

[3] F. J. Anscombe and R. J. Aumann, A definition of subjective probability, *Annals of Mathematical Statistics*, 34, 199–205, 1963.

[4] I. Aydogan, L. Berger, V. Bosetti and N. Liu, Three layers of uncertainty, *Journal of the European Economic Association*, 21, 2209–2236, 2023.

[5] R. Bansal and A. Yaron, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance*, 59, 1481–1509, 2004.

[6] A. Ben-Tal and M. Teboulle, An old-new concept of convex risk measures: The optimized certainty equivalent, *Mathematical Finance*, 17, 449–476, 2007.

[7] G. E. P. Box, Science and statistics, *Journal of the American Statistical Association*, 71, 791–799, 1976.

[8] R. Casadesus-Masanell, P. Klibanoff and E. Ozdenoren, Maxmin expected utility over Savage acts with a set of priors, *Journal of Economic Theory*, 92, 35–65, 2000.

[9] S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci and L. Montrucchio, Uncertainty averse preferences, *Journal of Economic Theory*, 146, 1275–1330, 2011.

[10] S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci and L. Montrucchio, Classical subjective expected utility, *Proceedings of the National Academy of Sciences*, 110, 6754–6759, 2013.

[11] G. Chamberlain, Robust decision theory and econometrics, *Annual Review of Economics*, 12, 239–271, 2020.

[12] P. Collin-Dufresne, M. Johannes and L. A. Lochstoer, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review*, 106, 664–698, 2016.

[13] D. R. Cox, Comment on "Model uncertainty, data mining and statistical inference", *Journal of the Royal Statistical Society*, Series A, 158, 455–456, 1995.

[14] T. Denti and L. Pomatto, Model and predictive uncertainty: A Foundation for smooth ambiguity preferences, *Econometrica*, 90, 551–584, 2022.

[15] I. Esponda and D. Pouzo, Berk-Nash equilibrium: A framework for modeling agents with misspecified models, *Econometrica*, 84, 1093–1130, 2016.

[16] D. Fudenberg, G. Romanyuk and P. Strack, Active learning with a misspecified prior, *Theoretical Economics*, 12, 1155–1189, 2017.

[17] T. Gajdos, T. Hayashi, J.-M. Tallon and J.-C. Vergnaud, Attitude toward imprecise information, *Journal of Economic Theory*, 140, 27–65, 2008.

[18] P. Ghirardato, F. Maccheroni, M. Marinacci and M. Siniscalchi, A subjective spin on roulette wheels, *Econometrica*, 71, 1897–1908, 2003.

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

[20] I. Gilboa, F. Maccheroni, M. Marinacci and D. Schmeidler, Objective and subjective rationality in a multiple prior model, *Econometrica*, 78, 755–770, 2010.

[21] I. Gilboa and D. Schmeidler, Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics*, 18, 141–153, 1989.

[22] L. P. Hansen, Beliefs, doubts and learning: Valuing macroeconomic risk, *American Economic Review*, 97, 1–30, 2007.

[23] L. P. Hansen, Nobel lecture: Uncertainty outside and inside economic models, *Journal of Political Economy*, 122, 945–987, 2014.

[24] L. P. Hansen and M. Marinacci, Ambiguity aversion and model misspecification: An economic perspective, *Statistical Science*, 31, 511–515, 2016.

[25] L. P. Hansen and J. Miao, Aversion to ambiguity and model misspecification in dynamic stochastic environments, *Proceedings of the National Academy of Sciences*, 115, 9163–9168, 2018.

[26] L. P. Hansen and T. J. Sargent, Robust control and model uncertainty, *American Economic Review*, 91, 60–66, 2001.

[27] L. P. Hansen and T. J. Sargent, Recursive robust estimation and control without commitment, *Journal of Economic Theory*, 136, 1–27, 2007.

[28] L. P. Hansen and T. J. Sargent, *Robustness*, Princeton University Press, Princeton, 2008.

[29] L. P. Hansen and T. J. Sargent, Macroeconomic uncertainty prices when beliefs are tenuous, *Journal of Econometrics*, 223, 222–250, 2021.

[30] L. P. Hansen and T. J. Sargent, Structured ambiguity and model misspecification, *Journal of Economic Theory*, 199, 2022.

[31] L. P. Hansen and T. J. Sargent, Risk, ambiguity, and misspecification: Decision theory, robust control, and statistics, *Journal of Applied Econometrics*, in press, 2023.

[32] G. H. Hardy, J. E. Littlewood and George Pólya, *Inequalities*, Cambridge University Press, Cambridge, 1952.

[33] E. Hewitt and L. J. Savage, Symmetric measures on Cartesian products, *Transactions of the American Mathematical Society*, 80, 470–501, 1955.

[34] D. H. Jacobson, Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games, *IEEE Transactions on Automatic Control*, AC-18, 124–131, 1973.

[35] P. Klibanoff, M. Marinacci and S. Mukerji, A smooth model of decision making under ambiguity, *Econometrica*, 73, 1849–1892, 2005.

[36] G. Lanzani, Dynamic concern for misspecification, mimeo, 2023.

[37] F. Liese and I. Vajda, *Convex statistical distances*, Teubner, Leipzig, 1987.

[38] F. Maccheroni, M. Marinacci and A. Rustichini, Ambiguity aversion, robustness, and the variational representation of preferences, *Econometrica*, 74, 1447–1498, 2006.

[39] M. Marinacci, Model uncertainty, *Journal of the European Economic Association*, 13, 1022–1100, 2015.

[40] L. Nascimento and G. Riella, A class of incomplete and ambiguity averse preferences, *Journal of Economic Theory*, 146, 728–750, 2011.

[41] I. R. Petersen, M. R. James and P. Dupuis, Minimax optimal control of stochastic uncertain systems with relative entropy constraints, *IEEE Transactions on Automatic Control*, 45, 398–412, 2000.

[42] D. Schmeidler, Subjective probability and expected utility without additivity, *Econometrica*, 57, 571–587, 1989.

[43] T. Strzalecki, Axiomatic foundations of multiplier preferences, *Econometrica*, 79, 47–73, 2011.

[44] A. Wald, *Statistical decision functions*, John Wiley & Sons, New York, 1950.

[45] S. G. Walker, Bayesian inference with misspecified models, *Journal of Statistical Planning and Inference*, 143, 1621–1633, 2013.