Robust Bayesian choice☆

Lorenzo Stanca

Collegio Carlo Alberto and University of Turin, Department of ESOMAS, Corso Unione Sovietica, 218 Bis, 10134, Turin, Italy

ARTICLE INFO

ABSTRACT

A major concern with Bayesian decision making under uncertainty is the use of a single probability measure to quantify all relevant uncertainty. This paper studies prior robustness as a form of continuity of the value of a decision problem. I show that this notion of robustness is characterized by a form of stable choice over a sequence of perturbed decision problems, in which the available acts are perturbed in a precise fashion. I then introduce a choice-based measure of prior robustness and apply it to models of climate mitigation and portfolio choice.

1. Introduction

The last 30 years have seen an extensive development of non-Bayesian theories of choice under uncertainty. Starting with the seminal work of Schmeidler (1989) and Gilboa and Schmeidler (1989), economists have developed models that depart from the standard subjective expected utility model. The study of such departures is motivated by different types of considerations.

For one, experimental evidence such as Ellsberg's paradox has suggested the Bayesian approach is not consistent with observed behavior. A second type of concern finds the Bayesian approach to be inadequate from a normative standpoint. As suggested by the literature on ambiguity and ambiguity aversion, a decision maker may find it challenging to specify a unique probability when only vague or fragmentary information is available.[1] Analogous concerns have emerged in other fields of economics.[2]

These concerns may be addressed by appealing to an informal continuity principle: even if the probabilities are not correctly specified, as long as the approximation error is small enough, then it should cause only a small variation in the final conclusions. This is a form of robustness to small specification errors of the original prior. For example, a key question within the literature on economy-climate models is to assess the robustness of the optimal policies to parameter assumptions in the specification of the damage function or in the temperature equation (e.g., see Ackerman et al. (2010)). In Bayesian

statistics, Kadane and Chuang (1978) formalize this idea by defining robustness as continuity of the value of a Bayesian decision problem under uncertainty.

In this paper, I provide a choice-theoretic foundation to this notion of robustness, i.e. I connect robustness to observable choice behavior.

To illustrate, consider the perspective of an analyst who observes choices over Savage acts made by an agent. The major difficulty with relating this type of robustness to choice behavior is that the analyst would have to be able to observe the agent's choices in the counterfactual scenario in which his belief is perturbed. This is not feasible in an observational study. Even in an experimental setting, reliably inducing perturbed beliefs may be challenging. Nonetheless, it is reasonable to assume that the analyst can change or perturb the acts available to the agent.

The approach I propose is to look at choice behavior over "perturbed" decision problems, i.e., decision problems in which the available acts are perturbed in a precise fashion. Following this reasoning, I provide a behavioral axiom that consists in stable (or convergent) choice behavior over a sequence of perturbed decision problems. The main result, Theorem 2, states that robustness is characterized by this form of stable choice. Therefore, one can think of robustness equivalently as a form of robust choice behavior over perturbed decision problems.

I then study how to quantify prior robustness by constructing a measure drawing from methods in functional differentiation. For an

---

[1] See Gilboa and Marinacci (2016) for a review of the literature on ambiguity aversion.

[2] A notable example is Hansen and Sargent's work in macroeconomics. In a series of influential papers (e.g., Hansen and Sargent (2001)) they considered decision makers who view their model (i.e., a probability distribution) as an approximation and want to behave robustly to possible perturbations of this approximating model.

agent with utility u, prior P and optimal act f^* robustness is quantified by

$$\sup_{Q \in C} \int u(f^*) dQ - \int u(f^*) dP, \tag{1}$$

where C is a set of probability measures that represent perturbations to the prior P. The second main set of results, Theorem 3 and Proposition 2, provides a foundation for such a measure.

This measure can be used to address two different types of questions. First, it can be used to assess how sensitive the predictions of a model are to the choice of the initial probability. For instance, it can be applied to Bayesian statistical methods to compare the robustness of different priors.

From a decision-theoretic perspective, this measure can be used to compare attitudes toward robustness for different agents. Consider two agents with the same utility but different beliefs. Proposition 2 shows that an agent is associated with a lower measure of robustness than another agent if and only if the monetary value he attaches to having his optimal act perturbed is lower than that of the other agent. In other words, a "more robust" agent will be less affected by perturbations of the optimal act.

I provide two applications to illustrate the importance of this measure of robustness: a climate mitigation problem and a portfolio choice problem.

An extensive literature (e.g., see Weitzman (2011) or Ibragimov et al. (2015)) has suggested that adopting "fat" or "heavy" tailed distributions is a way to build models with more robust conclusions. For example, fat tailed distributions such as the Student's t-distribution are typically considered a robust alternative to the use of normal distributions. I consider a simple climate mitigation model, where an agent has to choose the consumption of a good that can produce (an uncertain) damage in the future. A desire for robustness may emerge from experts' disagreement about the distribution of future damage.[3] I show that the measure of robustness I develop ranks as more robust distributions with heavier tails. One way to interpret this result is that with heavier tails, social utility will be less volatile to misspecification of the prior probability.

Further, in a simple portfolio allocation problem, I show that if the utility function incorporates explicitly a distaste for fat tails, modeling returns of a risky asset with a Student's t-distribution is more robust than it is with a normal distribution. Hence, this measure of robustness formalizes the intuition in the literature that connects heavy tailed distributions with robustness. As I discuss, heavy tails can be seen as emerging from model uncertainty, thus showing – against the common claim in the decision-theoretic literature – that the Bayesian approach can be properly used to deal with model uncertainty.

Related literature. Modeling robustness in a Bayesian framework is an old topic of interest. For instance, Savage et al. (1963) introduce the so called principle of "stable estimation". In a Bayesian statistical problem, they propose conditions such that the likelihood function dominates the prior distribution. Thus, robustness is modeled by the fact that the prior does not have a strong influence on the posterior. Fishburn et al. (1968) describe a variety of methods that may be used to evaluate the robustness of probabilities. The main approach they consider is to evaluate how much a probability that guarantees a unique optimal solution has to be perturbed to change the optimum. An extension of their work is given by Pierce and Folks (1969). Dempster (1975) contains a very interesting discussion of conceptual issues related to robustness from a subjectivist perspective.

In game theory, it is well known that game-theoretic predictions can be highly sensitive to assumptions about players' higher-order beliefs. Rubinstein's (1989) seminal paper shows that a strict Nash equilibrium of a game might fail to be rationalizable under a slight perturbation of hierarchies of beliefs. This paper spawned a large literature that tried to establish whether in some cases robustness can be preserved. From a theoretical perspective, this literature is close to my approach since it studies robustness of game-theoretic predictions with respect to a fixed topology. It is important to note that one could adopt an alternative approach suggested by Dekel et al. (2006) which constructs the coarsest metric topology that preserves a form of robustness.

This paper is inspired by several works in the literature in statistics. Good's writings on robustness (see for example Good, 1971) had a large impact on the literature in Bayesian statistics. Some of his insights were eventually developed in the literature on Bayesian robustness (see Berger et al. (2000) for a review). This literature studies how much Bayesian statistical methods depend on uncertainty about the precise details of the analysis, typically those given by the prior distribution. The main notion of robustness in this paper is inspired by this literature, and in particular by the work of Kadane and Chuang (1978).

This paper is also related to the frequentist literature on robustness that started with the seminal work of Huber et al. (1964). In particular, there are tight connections with the work of Hampel (1971, 1974). The main differences between my work and the literature in statistics is that I adopt a choice-theoretic approach.

Structure. Section 2 introduces the formal decision-theoretic framework. Section 3 introduces the notion of robustness, along with its behavioral characterization. Section 4 studies how to quantify the robustness for a given decision problem. Section 5 studies applications, and Section 6 offers concluding remarks. The proofs are in Appendix.

2. Preliminaries

Choice setting. I adopt the standard decision theoretic set-up à la Savage with additional assumptions on the state space and the set of consequences. The set S represents the states of the world and $\Sigma = \{A, E, \ldots\}$ is a σ-algebra of subsets of S called events. $X = \{x, y, z, \ldots\}$ is the set of consequences. Assume X is a Euclidean space, i.e. $X = \mathbb{R}^n$, and denote with $\| \cdot \|$ the standard Euclidean norm. $C_b(S, X) \subseteq X^S$ denotes the set of continuous and bounded functions. For example, if $S = [0, 1]$ then $C_b([0, 1], \mathbb{R}) \equiv C([0, 1])$ is the set of continuous real valued functions defined on the interval $[0, 1]$.,

I assume S is a Polish space and that Σ is the Borel σ-algebra. Δ denotes the set of countably additive probability measures $\mu : \Sigma \to [0, 1]$. Call $\mu \in \Delta$ *non-atomic* if for every $A \in \Sigma$ there is $B \in \Sigma$ such that $B \subseteq A$ and $\mu(A) > \mu(B) > 0$. With $ca(\Sigma)$ I denote the set of all countably additive signed measures defined on Σ. By standard results (see Dunford and Schwartz (1958)) since S is a Polish space, $ca(\Sigma)$ is isometrically isomorphic to the dual space of $C_b(S, \mathbb{R})$. One can therefore consider the duality pairing $\langle \cdot, \cdot \rangle : ca(\Sigma) \times C_b(S, \mathbb{R}) \to \mathbb{R}$ given by $\langle P, \phi \rangle = \int \phi \, dP$.

Given a sequence $(P_n)_n$ in Δ, $P_n \xrightarrow{*} P$ denotes convergence in the weak* topology, that is

$$\lim_{n \to \infty} \langle P_n, \phi \rangle = \langle P, \phi \rangle, \text{ for all } \phi \in C_b(S, \mathbb{R}).$$

I denote with $\mathcal{F} = \{f, g, h \ldots\} \subseteq X^S$ the set of *acts*. As usual, for $x \in X$, I define $x \in \mathcal{F}$ to be the constant act such that $x(s) = x$ for all $s \in S$. For any $f, g \in \mathcal{F}$ and event $A \in \Sigma$ denote with $f A g$ the act h such that $h(s) = f(s)$ for $s \in A$ and $h(s) = g(s)$ for $s \notin A$. A simple act is an act with finite support. Because X is a normed vector space, any simple act f can be written as $f = \sum_{i=1}^n 1_{A_i} x_i$, where 1_{A_i} is the indicator function of the set A_i, $(A_i)_{i=1}^n$ is a Σ-measurable partition of S and $x_1, \ldots, x_n$ are the elements of the range of f.

---

[3] In the case of climate change, there is substantial disagreement among experts on the parameter of climate sensitivity, i.e., by how much global average temperatures increase as a result of increased greenhouse gas levels. See Meinshausen et al. (2009).

The starting point of the analysis is a binary relation $\succcurlyeq$ on the set $\mathcal{F}$ that represents a decision maker's (DM) preferences over acts. Given acts f, g I write $f \succcurlyeq g$ if $f(s) \succcurlyeq g(s)$ for every $s \in S$. An act $f \in \mathcal{F}$ is *measurable* if $\{s \in S : f(s) \succcurlyeq x\} \in \Sigma$ and $\{s \in S : x \succcurlyeq f(s)\} \in \Sigma$ for every $x \in X$. I restrict attention to bounded and measurable acts. In other words,

$$\mathcal{F} = \{f \in X^S : f \text{ is measurable and } y \leqslant f \leqslant x \text{ for some } x, y \in X\}.$$

A functional $V : \mathcal{F} \to \mathbb{R}$ represents $\succcurlyeq$ if

$$V(f) \geq V(g) \iff f \succcurlyeq g,$$

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

For $g : S \to \mathbb{R}$ and a measure μ such that g is μ-integrable, let

$$\int g \, d\mu,$$

denote the standard Lebesgue integral with respect to μ. If ν, μ are two measures then $\int g \, d(\mu + \nu)$ denotes the integral with respect to the measure $A \mapsto \mu(A) + \nu(A)$ for all $A \in \Sigma$.

Given a measurable space $(\Omega, \mathcal{A})$, $\Delta(\Omega)$ denotes the set of all countably additive probability measures defined on $\mathcal{A}$. Given $P, Q \in \Delta(\Omega)$, I write $Q \ll P$ if and only if $A \in \mathcal{A}$, $P(A) = 0 \implies Q(A) = 0$. In some examples and applications, the integral of a function $g : \Omega \to \mathbb{R}$ with respect to a probability measure μ will be denoted by $\mathbb{E}_\mu g(\omega)$.

2.1. Basic preference representation

The DM's preference relation $\succcurlyeq$ over $\mathcal{F}$ is assumed to be represented by $V : \mathcal{F} \to \mathbb{R}$ satisfying

$$V(f) = \int u(f) \, dP \quad \forall f \in \mathcal{F}, \tag{2}$$

where $P \in \Delta$ is non-atomic and $u : X \to \mathbb{R}$ is continuous. One could provide a foundation for this criterion by means of Kopylov's (2010) characterization of Savage's subjective expected utility with countably additive probabilities, with the addition of a standard continuity axiom.

3. Robustness

Decision problems and robustness. A *decision problem* is a (non-empty) set $F \subseteq \mathcal{F}$ of acts. Acts in F are the available acts that the DM can choose. Since $\succcurlyeq$ satisfies Savage's axioms, the DM faces the usual optimization problem

$$\sup_{f \in F} \int u(f) \, dP. \tag{3}$$

Many economic models involve an optimization problem like (3). I will consider two main examples.

Example 1 (*Portfolio Choice*). In the standard portfolio choice problem, there are two assets available: a risk free one with certain return r_f and a risky one with uncertain return described by the random variable $r : \Omega \to \mathbb{R}$ defined on a measurable space $(\Omega, \mathcal{A})$. The investor has to allocate his wealth (which I normalize to 1) between the two assets. He cares about his terminal wealth $w \in \mathbb{R}$ and has utility $v(w)$. The set of available acts can be written as

$$A = \{ar + (1-a)r_f : a \in [0,1]\},$$

where a denotes the fraction of wealth invested in the risky asset. The problem faced by the investor is

$$\max_{r' \in A} \mathbb{E}_p v(r'(\omega)).$$

where $p \in \Delta(\Omega)$. Thus here we have $X = \mathbb{R}$, $S = \Omega$, $u = v$, $P = p$ and $F = A$.

Example 2 (*Climate Mitigation*). Consider a simple economic model of climate mitigation analogous to that studied in Gollier et al. (2000) (see also Bommier et al. (2021)). There are two periods where the only source of utility comes from the consumption of a good c_t at $t = 1, 2$. Consuming the good at $t = 1$ is free and certain but it may reduce the (uncertain) value of consumption at time $t = 2$ through environmental damage. More formally, the decision maker has to choose the level of climate abatement $a \in \mathbb{R}_+$, which will result in reduced consumption at $t = 1$, i.e. $c_1(a) = \bar{c} - r(a)$, where $r : \mathbb{R}_+ \to \mathbb{R}$ is a function that describes the cost of the abatement policy. At the same time, a higher level of abatement policy will (potentially) increase the future level of consumption $c_2(a, s)$ depending on the realization of a state $s \in \mathbb{R}$. The optimization problem is therefore given by:

$$V(P) = \max_{a \in \mathbb{R}_+} v(\bar{c} - r(a)) + \beta \mathbb{E}_P \left(c_2(a, \cdot) \right),$$

where $\beta \in (0, 1]$ reflects time preference and $P \in \Delta(\mathbb{R})$ is the decision maker belief about the state $s \in S$. The set of available acts can be written as

$$A = \{(x_1, x_2) : S \to \mathbb{R}^2 : x_1 = w_1 - c(a), x_2 = c_2(a, \cdot), a \in \mathbb{R}_+\}.$$

As the previous examples illustrate, it is common to make regularity assumptions on the set of feasible acts. To study robustness from a choice-theoretic perspective, I am going to make the following assumptions on F.

Assumption 1 (*Continuous Acts*). $F \subseteq C_b(S, X)$

Assumption 2 (*Optimal Act*). There exists $f^* \in F$ such that $f^* \succcurlyeq f$ for every $f \in F$.

A few comments are in order. Assumption 1 is a standard regularity assumption. While this assumption excludes simple acts, the latter can be arbitrarily approximated by the former.[4]

Using continuous acts substantially eases the exposition; however, it is possible to allow for non-continuous acts. Further, this assumption guarantees that we can endow F with the sup-norm topology, i.e., the distance defined by

$$\| f - g \|_\infty = \sup_{s \in S} \| f(s) - g(s) \|.$$

Assumption 2 simply states that the optimization problem is "interesting" in the sense that it admits a solution.

Two notions of robustness are studied in this paper, which are adaptions to the present setting of Kadane and Chuang's (1978) notions of stability (see also Salinetti (1994)). The first is a form of continuity of the value of the decision problem as a function of the DM's belief.

Definition 1. Fix a decision problem F and consider the preference $\succcurlyeq$ with representation (u, P). Say that $\succcurlyeq$ is *robust* with respect to F if for every sequence $(P_n)_{n=1}^\infty$ in Δ such that $P_n \xrightarrow{*} P$ it holds

$$\sup_{f \in F} \int u(f) \, dP_n \to \max_{f \in F} \int u(f) \, dP,$$

as $n \to \infty$.[5]

---

[4] Formally, this fact is known as Lusin's Theorem; see for example Theorem 12.8 in Aliprantis and Border (2006).

[5] It is important to observe that since for some x it holds $x \succcurlyeq f^*$, we obtain $x \succcurlyeq f$ for every $f \in F$. This implies that

$$\sup_{f \in F} \int u(f) \, dP_n \leq u(x) < \infty,$$

for every n. Thus the sequence $(\sup_{f \in F} \int u(f) \, dP_n)_{n=1}^\infty$ is effectively a sequence in $\mathbb{R}$.

Observe Robustness does not depend on the normalization of the Bernoulli utility u. It follows that robustness is a property of the preference $\succcurlyeq$ which holds or does not hold for every expected utility representation. The next example provides a case in which robustness is not satisfied.

Example 3. Suppose that $\Omega = \mathbb{R}$ endowed with the Borel σ-algebra. Assume that $X = \mathbb{R}$, $u(x) = -x^2$ for every $x \in \mathbb{R}$, and

$$F = \{f : \mathbb{R} \to \mathbb{R} : f(\omega) = a - \omega \text{ for some } a \in \mathbb{R}\}.$$

Further assume that P is such that $\int \omega dP < \infty$ and $\int \omega^2 dP < \infty$. Under these assumptions, this decision problem has finite value.

For any positive integer n, let P_n denote perturbation of the original belief P given by $P_n = (1 - \frac{1}{n})\mu + \frac{1}{n}\delta_{n^2}$, where δ_{n^2} denotes the degenerate distribution that assigns probability one to n^2. Note that the sequence of distribution functions of $(P_n)_n$ converges pointwise to that of P. It follows that $P_n \overset{*}{\to} P$.

Now observe that the act f_n^* that maximizes $\int u(f)dP_n$ satisfies for every $\omega \in \Omega$

$$f_n^*(\omega) + \omega = \int \omega dP_n = \left(1 - \frac{1}{n}\right) \int \omega dP + \frac{1}{n}n^2 = \left(1 - \frac{1}{n}\right) \int \omega dP + n,$$

so that $f_n(\omega) \to \infty$ for every $\omega \in \Omega$. It follows that $\sup_{f \in F} \int u(f)dP_n \to -\infty$. In words, a very small perturbation of the initial belief leads the agent to extremely different conclusions.

The main reason I consider robustness with respect to the weak* topology is that it is a natural topology to consider for an initial analysis. Indeed, as the name suggests, it is weaker than most other topologies on probability measures that are typically considered. As a consequence, this notion of robustness will be very demanding. In particular, it will be more demanding than any other notion of robustness that uses a topology stronger than the weak*.

An important feature of this notion of robustness is that it depends not only on the prior probability but also on the choice set F and the utility u. As will be shown in the main theorem of this section, robustness is characterized by a form of robust choice behavior over sequences of perturbed decision problems.

Therefore, since more regular decision problems are harder to perturb, regularity properties of the decision problem will be important to guarantee robustness. For example, a key property related to robustness is the compactness of the choice set F.

Example 4 (*Portfolio Choice Continued*). Suppose that Ω is a metric space and $r : \Omega \to \mathbb{R}$ is a continuous bounded function. For example, $\Omega = [0, 1]$, $r(\omega) = \omega$ and P has a beta distribution (so that r is also distributed as a beta). Then it is easy to show that the set

$$A = \{ar + (1 - a)r_f : a \in [0, 1]\} \subseteq C([0, 1]),$$

is a compact subset of $C([0, 1])$ in the sup-norm topology by a direct application of the Arzelà–Ascoli theorem. Suppose that $v : \mathbb{R} \to \mathbb{R}$ is continuous. Given the structure of the feasible set A, it is possible to show that the objective function is continuous in both the probability and the choice variable.

Thus, by the maximum theorem (see for example p. 306 in Ok (2011)) robustness is satisfied. On the other hand, in general strong robustness will not be satisfied unless there is a unique optimal solution. For the general decision problem in (3), the following result holds.

Proposition 1. *Suppose that F is a compact subset of $C_b(S, X)$. Then robustness is satisfied. Moreover, strong robustness is satisfied whenever there is a unique optimal act.*

Proof. See Appendix. $\quad\square$

Observe that the previous result does not rely on the maximum theorem. This is due to the fact that the objective function need not be jointly continuous in both its arguments. However, part of the proof of Theorem 2 can be used to show that in this example robustness is satisfied. The full proof is elaborated in Appendix.

In general, for a decision problem $F' \subseteq \mathcal{F}$, any constant $\varepsilon > 0$ and any expected utility representation (u, P) of preferences, we will be interested in the set of all ε-optimal acts

$$\max_{u, P, \varepsilon}(F') = \left\{ f \in F' : \int u(f)dP \geq \sup_{g \in F'} \int u(g)dP - \varepsilon \right\},$$

while the set of optimal acts is denoted with $\max_{u, P}(F')$.[6]

The second notion of robustness I look at requires a form of continuity of the optimal solution.

Definition 2. Fix a decision problem F and consider the preference $\succcurlyeq$ with representation (u, P). Say that $\succcurlyeq$ is *strongly robust* with respect to F if for every sequence $(P_n)_{n=1}^{\infty}$ in Δ such that $P_n \overset{*}{\to} P$ and for every sequence of positive numbers $(\varepsilon_n)_{n=1}^{\infty}$ such that $\varepsilon_n \to 0$, every sequence of acts $(f_n)_{n=1}^{\infty}$ that satisfies $f_n \in \max_{u, P_n, \varepsilon_n}(F)$ converges to an optimal act $f^* \in \max_{u, P}(F)$.

Remark 1. Strong robustness is indeed stronger than robustness. See Theorem 4 in Appendix.

3.1. Perturbed decision problems

The main question that I study is whether it is possible to obtain a behavioral characterization of these two notions of robustness. The challenge with this question is that it requires observing the choices made by the DM under alternative beliefs. Unfortunately, in many settings such a counterfactual is not available. The main idea that I propose is that one can look at the behavior of the DM when the decision problem itself is "perturbed" in a precise fashion.

Given an act f, I will consider "perturbations" of the kind fEx for some event $E \in \Sigma$ and outcome $x \in X$. More precisely, given a sequence $(P_n)_{n=1}^{\infty}$ such that $P_n \overset{*}{\to} P$, consider the decision problem F_n defined by

$$F_n = \left\{ fEx : f \in F, E \in \Sigma, x \in X, \int u(fEx)dP = \int u(f)dP_n \right\}.$$

F_n contains all the perturbations of the acts in F that have expected utility "as if" the agent's belief was P_n. Because the sets $(F_n)_n$ contain perturbations of the acts in F, they will not necessarily satisfy Assumptions 1 and 2.[7]

The first result describes an important class of such perturbations and, in particular, shows that $F_n \neq \varnothing$ for every n.

Theorem 1. *Let F be a decision problem and $(P_n)_n$ a sequence such that $P_n \overset{*}{\to} P$. For every $f \in F$, there exist $A_{f,n}$ and $x_{f,n}$ such that*

$$\int u(f A_{f,n} x_{f,n})dP = \int u(f)dP_n \quad \forall n.$$

Moreover, $P(A_{f,n}) \to 1$ and either $x_{f,n} = x_f$ or $x_{f,n} = y_f$, where $x_f \succcurlyeq f \succcurlyeq y_f$.

Proof. See Appendix. $\quad\square$

---

[6] Observe that the set $\max_{u, P, \varepsilon}(F')$ is non-empty for every $\varepsilon > 0$. Indeed, because $\sup_{g \in F'} \int u(g)dP < \infty$, there exists a sequence $(f_n)_n$ in F' such that $\lim \int u(f_n)dP = \sup_{g \in F'} \int u(g)dP$.

[7] Because an act $g \in F_n$ will have the form $g = fEx$, in general it will not satisfy continuity, that is $g \notin C_b(S, X)$.

Thus, not only are the sets F_n non-empty but for each $f \in F$ there is some $fEx \in F_n$ that, for large n, is "close" to the act f. Indeed, it is easy to show that

$$\int \| f A_{f,n} x_{f,n} - f \| dP \to 0 \text{ as } n \to \infty.$$

In this sense, the perturbed act $f A_{f,n} x_{f,n}$ can be considered a small alteration of the act f for large n.

Thanks to this result, it is possible to understand the choice behavior of the DM in the counterfactual scenario in which he had a different belief. Thus, given any $P_n \overset{*}{\to} P$, it is possible to understand the DM's behavior as if his belief was P_n by looking at choices over the set F_n.

The next key property captures the idea of *stable* (or convergent) choice over the sequence of perturbed decision problems $(F_n)_{n=1}^\infty$.

Definition 3. Consider $\succeq$ with representation (u, P) and a decision problem $F \subseteq \mathcal{F}$. Let $\varepsilon_n \to 0$. A sequence $(g_n)_{n=1}^\infty = (f_n E_n x_n)_{n=1}^\infty \in \prod_{n=1}^\infty \max_{u,P,\varepsilon_n}(F_n)$ is stable for the optimal act $f^* \in \max_{u,P}(F)$ if the following two conditions hold:

(i) There is a subsequence $(f_{n_k})_k$ of $(f_n)_{n=1}^\infty$ such that $f_{n_k} \to f^*$;

(ii) $\int \| g_{n_k} - f^* \| dP \to 0$.

Stability requires a strong type of convergence for the sequence

$(g_{n_k})_k = (f_{n_k} E_{n_k} x_{n_k})_k$.

First, the sequence of acts $(f_{n_k})_k$ that are perturbed has to convergence to an optimal act f^* according to the sup-norm metric. Moreover, the sequence $(g_{n_k})_k = (f_{n_k} E_{n_k} x_{n_k})_k$ has to converge to the optimal act f^*.

To have intuition for condition (ii), note it requires the usual convergence in mean of the sequence $(g_{n_k})_k$ to f^*. Thus, choice behavior is stable in the sense that choices for the perturbed decision problems are similar to *some* choices potentially observed in the original problem.

The main axiom of this paper requires stable behavior over sequences of perturbed decision problems.

Axiom (*Preference for Stability*). Consider $\succeq$ with representation (u, P). Fix a decision problem $F \subseteq \mathcal{F}$. Say that $\succeq$ has a preference for stability with respect to F if for every $P_n \overset{*}{\to} P$ there exists $\varepsilon_n \to 0$ and a stable sequence $(g_n)_{n=1}^\infty = (f_n E_n x_n)_{n=1}^\infty \in \max_{u,P,\varepsilon_n}(F_n)$ for some optimal act f^*.

In words, preference for stability requires stable choice behavior over any sequence of perturbed decision problems. The main result of the paper characterizes robustness in terms of preference for stability.

Theorem 2. $\succeq$ *is robust with respect to F if and only if it has preference for stability with respect to F.*

Proof. See Appendix. $\square$

As a corollary, an analogous characterization can be obtained for strong robustness.

Corollary 1. $\succeq$ *is strongly robust if and only if for $\varepsilon_n \to 0$ and every sequence $(f_n)_{n=1}^\infty$ such that $f_n \in \max_{u,P_n,\varepsilon_n}(F)$ there exists a sequence $(g_n)_{n=1}^\infty = (f_n E_n x_n)_{n=1}^\infty$ such that $g_n \in \max_{u,P,\varepsilon_n}(F_n)$ and for a subsequence $(g_{n_k})_{k=1}^\infty$ satisfies*

$$\int \| g_{n_k} - f^* \| dP \to 0.$$

Conceptually, these two results connect the notions in Definition 1 to a property of choice that is potentially testable. Clearly, an axiom involving convergence of choices is not directly "operational", since it involves an infinite sequence of acts.

Nonetheless, these results do suggest that attitudes toward robustness are related to convergence of choices under small perturbations of the available acts. The fact that the axiom involves convergence of a sequence of acts does not necessarily preclude testability. For

instance, the experimental literature that studies learning in games (e.g., Hyndman et al. (2012)) has studied convergence of actions in situations of repeated interaction.

Typically, convergence is assumed whenever the same type of choice is repeatedly observed over a period of time. Therefore, one could understand whether robustness fails by looking at whether or not choices over a sequence of perturbed decision problems converge to the choice for the original problem. To clarify this point, consider the following example.

Example 5. Consider the following special case of Example 1. Here $\Omega = [0, 1]$, P is the Lebesgue measure, and $u(x) = \log(x)$. Assume that the riskless asset pays $\frac{1}{3}$ for sure. Given these assumptions, the optimal allocation of wealth is to allocate $\alpha \approx 71.63\%$ of wealth to the risky asset. Consider the perturbation of the prior P given by $P_n = \frac{1}{n}\delta_1 + (1 - \frac{1}{n})P$, where δ_1 is the distribution such that $\delta_1(\{1\}) = 1$. In words, this sequence of perturbations takes into account the possibility that the risky asset might involve no risk and pays 1 for sure. For every allocation of wealth α, it is possible to find a perturbation r_n of the risky asset given by $r_n(\omega) = \omega E_n 1$, where $E_n = [0, x_n(\alpha)] \subseteq \Omega$ and $x_n(\alpha) \to 1$ as $n \to \infty$ such that[8]

$$\mathbb{E}_P \log\left(\alpha \omega E_n 1 + (1 - a)\frac{1}{3}\right) = \mathbb{E}_{P_n} \log\left(\alpha \omega + (1 - a)\frac{1}{3}\right).$$

Therefore, for every n we can consider the perturbed decision problems

$$\max_{\alpha \in [0,1]} \mathbb{E}_P \log\left(\alpha \omega E_n 1 + (1 - a)\frac{1}{3}\right). \tag{4}$$

In these decision problems, the risky asset has an uncertain outcome for $\omega \in E_n$ (where E_n is close to Ω for large n) and pays 1 for sure otherwise. A failure of preference for stability would consist in observing choices of the allocation of wealth α for the problem (4) that for "large" n are very different from 71.63%.

4. A quantitative measure of robustness

This section develops a quantitative measure of robustness. Quantifying robustness is interesting for two main reasons. First, one may want to quantify the sensitivity of the predictions of a model to the choice of prior. Moreover, one may be interested in comparing attitudes toward robustness for different agents from a decision theoretic perspective. The measure is inspired by the concept of the influence function, which was developed by Hampel (1971, 1974) in the context of robust (frequentist) statistics.

Robustness will be measured with respect to some class of perturbations $C \subseteq \Delta$ of probability measures such that $P \in C$. This set can be interpreted as a set of perturbations considered plausible by the DM. At the end of this section, I offer a few possible specifications for the set C.

Define the map $W : \Delta \to \mathbb{R}$ by

$$W(\mu) = \sup_{f \in F} \int u(f) d\mu \text{ for every } \mu \in \Delta.$$

Given $Q \in C$, consider the affine directional derivative of W at $P \in \Delta$ in the direction Q

$$d_Q W(P) = \lim_{h \downarrow 0} \frac{W(hQ + (1 - h)P) - W(P)}{h}.$$

The affine directional derivative is equivalent to a standard directional derivative in the direction $Q - P$ at the point P. It is an intuitive way to

---

[8] To see this, note that $\mathbb{E}_P \log(\alpha \omega E_n 1 + (1 - a)) = \int_{x_n}^1 \log(\alpha \omega + (1 - \alpha)\frac{1}{3}) d\omega + (1 - x_n) \log(\alpha + (1 - \alpha)\frac{1}{3})$ and $\mathbb{E}_{P_n} \log(\alpha \omega + (1 - a)) = (1 - \frac{1}{n}) \int_0^1 \log(\alpha \omega + (1 - \alpha)\frac{1}{3}) + (\frac{1}{n}) \log(\alpha + (1 - \alpha)\frac{1}{3})$. Thus, one can pick x_n with $x_n \to 1$ so that the equality holds.

describe how the value changes when the prior P is perturbed by the probability Q.[9]

The main result of this section is an envelope theorem that gives an explicit formula for the directional derivative $d_Q W(P)$. Furthermore, it also provides a connection between robustness and differentiation of W.

Theorem 3. *Consider $\succeq$ with representation (u, P) and a decision problem $F \subseteq \mathcal{F}$ such that robustness is satisfied and there is a unique optimal act f^*. Suppose that X is a Euclidean space. Then it holds that*

$$d_Q W(P) = \int u(f^*) dQ - \int u(f^*) dP.$$

Proof. See Appendix. □

Therefore, under robustness the effect of an infinitesimal perturbation can be computed by simply comparing the value of the decision problem to the expected utility obtained by choosing the optimal act for P when the "true" probability is Q. Based on this result, for a given triple (u, P, F), robustness can thus be quantified by taking the $Q \in C$ that maximizes this difference.

Definition 4. For $\succeq$ with representation (u, P) and a decision problem F, define $m(u, P, F)$ as

$$m(u, P, F) = \sup_{Q \in C} |d_Q W(P)| = \sup_{Q \in C} \int u(f^*) dQ - \int u(f^*) dP. \quad (5)$$

This approach is similar to the one used in robust statistics (see Hampel (1974), pp. 387–388). It is straightforward to find conditions that guarantee that the supremum in (5) is attained.

Proposition 2. *If C is weak**-compact then there exists Q^* such that*

$$m(u, P, F) = \int u(f^*) dQ^* - \int u(f^*) dP.$$

Proof. Omitted. □

It is important to note that the magnitude of $\sup_{Q \in C} \int u(f^*) dQ - \int u(f^*) dP$ does not in itself have any meaning in utility theory, since it can be made arbitrarily large or small by an affine positive transformation of the utility function u. It can nonetheless be used to compare attitudes toward robustness of different agents who share the same Bernoulli utility and set of plausible perturbations C, as discussed in the next result. In applications, the primary interest lies in understanding how varying beliefs influence robustness, rather than uncertainties related to tastes.

Consider two preferences $\succeq_1$ and $\succeq_2$ with representation given by (u, P_1), (u, P_2) and fix a decision problem $F \subseteq \mathcal{F}$. Assume that $\succeq_1$ and $\succeq_2$ have well-defined certainty equivalents, i.e., for $i = 1, 2$ and every $f \in \mathcal{F}$ there exists $CE_i(f) \in X$ such that $CE_i(f) \sim_i f$.[10] Denote with f_i an optimal act for agent i and suppose that $CE_1(f_1) \sim_1 CE_2(f_2)$ for some $x^* \in X$, or equivalently that $\int u(f_1) dP_1 = \int u(f_2) dP_2$. This last assumption requires the two agents to be comparable in the sense that they assign the same value to the decision problem. For a given agent $i = 1, 2$ and $Q \in C$, define analogously to the previous section the set of perturbations of the optimal act f_i as

$$F_Q^i = \{ f_i E x : E \in \Sigma, x \in X, \int u(f_i E x) dP_i = \int u(f_i) dQ \},$$

and

$$F_C^i = \bigcup_{Q \in C} F_Q^i.$$

Note that by Theorem 1 the sets F_Q^i are non-empty. Denote with $f_i E_i x_i$ an optimal act in F_C^i for agent i. The next proposition will show that these are well-defined and will provide an interpretation for the statement that agent 2 is "more robust" than agent 1, i.e. $m(u, P_1, F) \geq m(u, P_2, F)$.

Proposition 3. *Assume that $X = \mathbb{R}$ and that the set C is weak**-compact. Then*

$$m(u, P_1, F) \geq m(u, P_2, F) \iff CE_1(f_1 E_1 x_1) \geq CE_2(f_2 E_2 x_2).$$

Proof. See Appendix. □

The interpretation of the previous result is simple: a more robust agent will value less (in monetary terms) the set of perturbations of the optimal act than the less robust agent.

Observe that since the ratio

$$\frac{m(u, P_1, F)}{m(u, P_2, F)} = \frac{\sup_{Q \in C} \int u(f^*) dQ - \int u(f^*) dP_1}{\sup_{Q \in C} \int u(f^*) dQ - \int u(f^*) dP_2},$$

is preserved under positive affine transformations of the utility function u, it can be used to measure the robustness of P_2 compared to that of P_1. In some cases, the value of $m(u, P, F)$ can be expressed analytically.

Example 6. Suppose that C is the Kullback–Leibler (KL) neighborhood used by Hansen and Sargent (see Strzalecki (2011)). More precisely, let

$$C = \{ Q \in \Delta : Q \ll P, R(Q \| P) \leq K \}, \quad (6)$$

where $K > 0$ and

$$R(Q \| P) = \begin{cases} \int \log \left(\frac{dP}{dQ} \right) dP & \text{if } Q \ll P; \\ \infty & \text{otherwise.} \end{cases}$$

The advantage of the KL neighborhood is that it is a very tractable non-parametric set of probability measures. Using well-known results, one can obtain a closed form representation for $m(u, P, F)$ when C is given by (6).

Proposition 4. *Suppose that C is given by (6). Then there exists $\theta \geq 0$ decreasing with K, such that*

$$m(u, P, F) = \theta \log \left(\int e^{\frac{1}{\theta} u(f^*)} dP \right) - \int u(f^*) dP.$$

Proof. See Appendix. □

Observe that in this example C need not be weak**-compact. Indeed, the comparative statics result in Proposition 3 will hold as long as the optimization problem has a solution, an assumption satisfied by the KL neighborhood.

5. Applications

This section provides applications of the measure of robustness to a climate mitigation problem (Example 1) and to a portfolio choice problem (Example 2). I show how under certain assumptions prior distributions with heavy or fat tails can be associated with higher robustness, as measured by the criterion developed in the previous section.[11]

Part of the literature on decision theory has suggested that the Bayesian framework cannot properly account for model uncertainty (see Marinacci (2015) and Berger and Marinacci (2020) pp. 487–489). It is argued that all uncertainty about the right probability or "model"

---

[9] A complete study of affine derivatives can be found in Cerreia-Vioglio et al. (2019).

[10] Certainty equivalents exist under standard regularity assumptions on u.

[11] I use the terms heavy and fat tails interchangeably, however some authors distinguish between the two; see for example Taleb and Cirillo (2019) p.6.

is reduced to risk: heuristically, given a set $(P_\theta)_\theta$ and a prior belief μ on θ the two-stage robust criterion problem

$$\sup_{f \in F} \int u(f) dP = \sup_{f \in F} \int \left(\int u(f) dP_\theta \right) d\mu,$$

is equivalent to the standard Bayesian criterion with average prior $\bar{P} = \int P_\theta d\mu(\theta)$.

However, the point here is that averaging different models can lead to thicker tails, e.g., the Student's t-distribution can be written as a mixture of Gaussian distributions.[12] Hence, model uncertainty can indeed lead to more robust decisions.

5.1. Climate mitigation

Consider the abatement policy from Example 2. Assume that $\beta = 1$, $v(x) = -e^{-x}$, $r(a) = a$, and that P is such that $c_2(a, \cdot) \sim \log \mathcal{N}(a\mu, \sigma^2)$ for every $a \in [0, \infty)$. Therefore, a higher level of abatement policy at $t = 1$ increases the average level of consumption at $t = 2$. The maximization problem can be written as:

$$V(P) = \max_{a \in \mathbb{R}_+} v\left(w_1 - r(a)\right) + \beta \mathbb{E}_P v\left(c_2(a, \cdot)\right) = \max_{a \in [0, \infty)} -e^{\frac{\sigma^2}{2} - a\mu} - e^{a - \bar{c}}.$$

Let

$$C = \{Q \in \Delta : Q \sim \log \mathcal{N}(\mu, \sigma), (\mu, \sigma) \in [\underline{\mu}, \bar{\mu}] \times [\underline{\sigma}, \bar{\sigma}]\}.$$

The set C can be thought of as a group of experts who vary in terms of the level of the parameter (μ, σ^2). Weitzman (2011) highlighted importance of heavy-tailed distribution to model robustness with respect to catastrophic outcomes. Here I use the parameter of kurtosis to measure how heavy tails are (see for example Müller et al. (1998)).

Hence, a higher level of variance σ^2 is associated with heavier tails. The next proposition offers a comparative statics result that shows how heavier tails are ranked as more robust. Each P can be identified with the pair of parameters (μ, σ^2) and denote with a_P the optimal action for the belief P.

Proposition 5. *Consider $P_1, P_2 \in C$ such that $\sigma_2^2 > \sigma_1^2$, $V(P_1) = V(P_2)$ and*

$$\frac{\sigma_2^2 - \sigma_1^2}{2} \leq \left(a_{P_2} \mu_2 - a_{P_1} \mu_1 \right) - (a_{P_2} - a_{P_1}) \bar{\mu}. \tag{7}$$

Then it holds that

$$m(v, P_1, A) \geq m(v, P_2, A).$$

Proof. See Appendix. □

The above proposition formalizes the relationship between heavy-tailed distributions and robustness of mitigation policy. Heavy tails lead to more robust choices, provided that the difference between σ_2^2 and σ_1^2 is small enough as described by the bound in (7). Several papers have employed heavy tailed distributions (e.g., Ikefuji et al. (2020) and Ackerman et al. (2010)) to model catastrophic climate risk.

Proposition 5 demonstrates that climate mitigation policies that are based on heavy-tailed distributions will be more robust to model misspecification. This means that even if the underlying assumptions are incorrect, the impact on social utility will be less severe. Further, it shows how the measure of robustness developed here can be applied.

5.2. Portfolio choice

The Student's t-distribution is often regarded as a robust alternative to the normal distribution. In statistics, it is typically employed in the

Table 1
Measure of robustness for different neighborhood sizes.

	$\theta = 30$	$\theta = 20$	$\theta = 2$
$m(u, F, P_1)$	0.16574	0.16690	0.1918
$m(u, F, P_2)$	0.16544	0.16650	0.1905
$\frac{m(u, F, P_1)}{m(u, F, P_2)}$	1.00181	1.00240	1.00682

rejection of outliers, as first pointed out in a paper by De Finetti (1961). For example, Meinhold and Singpurwalla (1989) study a robustification of the Kalman filter using multivariate Student's t-distributions. In economics, Weitzman (2007) studies implications to asset pricing of parameter (or model) uncertainty, which leads to fat tailed distributions.

Here I use the measure developed in the previous section to compare the robustness of the t-distribution to the normal distribution. More precisely, in a simple portfolio allocation problem, I show that if the utility function incorporates explicitly a distaste for fat tails, modeling returns of a risky asset with a Student's t-distribution is ranked as more robust than normally distributed returns.

Consider again Example 1, in which the DM has to allocate his wealth (normalized to 1) between a risk-free asset with return fixed at $r_f = 1$ and a risky asset with values in $\mathbb{R}$, so that the problem can be written as

$$\max_{\alpha \in [0,1]} \mathbb{E}_P u(\alpha \omega + (1 - \alpha)).$$

To compare different probabilistic assumptions on P, it is necessary to specify a set C of possible perturbations. A tractable specification that I adopt is the entropy neighborhood from Eq. (6). In particular, I compare $P_1, P_2 \in C$ such that $P_1 \sim N(0, 1.716)$, $P_2 \sim t(5)$ and take P_1 to be the center of the neighborhood, so that $C = \{Q \in \Delta(\mathbb{R}) : R(Q \| P_1) \leq K\}$.

To incorporate an explicit distaste for fat tails, I assume that $u(x) = x - x^4$. To understand such an assumption, note that the DM prefers higher expected value and dislikes higher fourth moment. Since kurtosis, a measure of heavy tails, is identified with the fourth moment, such a specification for u can be thought of as a way to model a distaste for fat tails.

Under these assumptions, the optimal allocations of wealth for the two probabilities are $\alpha_1 \approx 0.120665$, $\alpha_2 \approx 0.11752$, and the optimal values are the same, $\mathbb{E}_{P_1} u(\alpha_1 \omega + (1 - \alpha_1)) = \mathbb{E}_{P_2} u(\alpha_2 \omega + (1 - \alpha_2)) \approx 0.163662$ so that the one can use the comparative robustness result from Proposition 3.

The next result provides a closed-form expression for $m(u, P_i, F)$ analogous to that in Proposition 4.

Proposition 6. *There exists $\theta \geq 0$ decreasing with K such that for $i = 1, 2$*

$$m(u, P_i, F) = \theta \log \left(\mathbb{E}_P e^{\frac{1}{\theta} u(\alpha_i \omega + (1 - \alpha_i))} \right) - \mathbb{E}_{P_i} u\left(\alpha_i \omega + (1 - \alpha_i) \right).$$

Proof. See Appendix. □

Thanks to this result, it is possible to compare the values of $m(u, P_i, F)$, $i = 1, 2$ for different values of θ.[13] Recall that θ is decreasing with the size of the neighborhood K. Approximate values are reported in Table 1.

Since $m(u, F, P_1) > m(u, F, P_2)$ for all values of θ, P_2 is ranked as more robust than P_2. The main force behind this result is that the share of wealth invested in the risky asset is higher under the normal distribution (i.e., $\alpha_1 > \alpha_2$), implying that

$$\sup_{Q \in C} \mathbb{E}_Q u(\alpha_1 \omega + (1 - \alpha_1)) > \sup_{Q \in C} \mathbb{E}_Q u(\alpha_2 \omega + (1 - \alpha_2)).$$

---

[12] See for example Murphy (2012), p. 361.

---

[13] It is possible to check that $P_2 \in C$ for all the corresponding values of $\theta = 30, 20, 2$.

Moreover, the ratio $\frac{m(u,F,P_1)}{m(u,F,P_2)}$ increases as θ decreases, i.e. as the size of C increases the relative robustness of P_2 with respect to P_1 increases.

6. Concluding remarks

This paper examined the following question: can one develop a choice-based theory of robustness in a purely Bayesian framework? The main motivation has been the practical appeal of the Bayesian approach, irrespective of its ability to rationalize actual behavior under uncertainty.

I have provided a positive answer to the above question. The starting point of this theory is an axiomatization of Bayesian decision makers whose optimal value for a fixed decision problem is continuous in the prior.

The axiomatic approach is one of the major novelties of the paper. In conclusion, this paper presents a theory of comparative robustness that enables a formal comparison of the sensitivity of the conclusions of a Bayesian model to variations in the prior. This contribution allows for a more thorough evaluation of the robustness of Bayesian models.

Data availability

No data was used for the research described in the article.

Appendix

Mathematical preliminaries

Topological preliminaries. Let (T, τ) be a first-countable topological space (so that only sequences need to be considered). Given a sequence $(t_n)_{n=1}^{\infty}$ we denote convergence to a point $t \in T$ by $t_n \xrightarrow{\tau} t$. A double-indexed sequence is a mapping $t : \mathbb{N} \times \mathbb{N} \to T$.

Lemma 1. *Consider a double-indexed sequence* $(t_{n,m})_{(n \in \mathbb{N}, m \in \mathbb{N})}$ *such that*

(1) For every m, $t_{(n,m)} \xrightarrow{\tau} t_m$ for some $t_m \in T$.
(2) $t_m \xrightarrow{\tau} t$ for some $t \in T$.
Then there exists an increasing mapping $\iota : \mathbb{N} \to \mathbb{N}$ and with $\lim_{m\to\infty} \iota(m) = \infty$ such that

$$t_{n,\iota(n)} \xrightarrow{\tau} t.$$

Proof. See Attouch (1984), Corollary 1.18. □

As discussed in the main text, one of the main mathematical techniques for studying robustness is that of Γ-convergence. Γ-convergence is a notion of convergence for functionals germane to studying the convergence of optima and maximizers. Its usual formulation is for minimization problems. Here I present the analogous notion for maximization problems.

Definition 5. Let T be a first-countable topological space. A sequence of functions $F_n : T \to \mathbb{R}$ Γ-converges to a function $F : T \to \mathbb{R}$ if

(i) For every sequence $t_n \xrightarrow{\tau} t$,

$$F(t) \geq \limsup_{n\to\infty} F_n(t_n).$$

(ii) For every $t \in T$, there exists a sequence $t_n \xrightarrow{\tau} t$ such that

$$F(t) \leq \liminf_{n\to\infty} F_n(t_n).$$

If F_n Γ-converges to F, then I write

$$\Gamma\text{-}\lim F_n = F.$$

The assumption that T is first-countable is necessary to focus only on sequences and avoid the use of nets. Γ-convergence is tightly connected to perturbations of optimization problems as the next result shows. This result will be extremely important in the proof of Theorem 2.

Theorem 4 (*Attouch, 1984, Theorem 1.10*). *Consider a first-countable topological space T and a functional $F : T \to \mathbb{R}$. A sequence of functionals $F_n : T \to \mathbb{R}$ such that $\Gamma\text{-}\lim F_n = F$ and $\operatorname{argmax} F \neq \emptyset$ satisfies $\sup F_n \to \max F$ if and only if there exists $\varepsilon_n \to 0$ and a compact sequence $(t_n)_n$ such that t_n is ε_n-optimal for F_n.*

The following simple example shows how Γ-convergence is not enough to get convergence of suprema and also shows the key role played by compactness.

Example 7. Consider the sequence of functions $F_n : \mathbb{R} \to \mathbb{R}$ defined by

$$F_n(t) = \begin{cases} 1 & t \geq n, \\ \frac{t}{n} & 0 \leq t < n, \\ t & t < 0. \end{cases}$$

It is possible to show that $\Gamma\text{-}\lim F_n = F$ where F is defined by

$$F(t) = \begin{cases} 0 & t \geq 0, \\ t & t < 0. \end{cases}$$

This follows by applying Proposition 5.2 in Dal Maso (1993) and the fact that F_n converges to F uniformly on every bounded set. However, note that $\max_{t\in\mathbb{R}} F_n(t) = 1 \to 1 \neq \max_{t\in\mathbb{R}} F(x) = 0$ and $\operatorname{argmax}_{t\in\mathbb{R}} F_n(t) = n \to \infty$.

The main references for the literature on Γ-convergence are Attouch (1984), Dal Maso (1993) and Braides (2002). An important notion of convergence related to Γ-convergence is that of Kuratowski convergence. Given a sequence $(C_n)_{n=1}^{\infty}$ of subsets of T, let

$$\operatorname{Ls} C_n = \{t \in T : \text{ there exist } (n_k)_{k=1}^{\infty} \text{ and } t_{n_k} \in C_{n_k} \text{ such that } t_{n_k} \xrightarrow{\tau} t\}.$$

and

$$\operatorname{Li} C_n = \{t \in T : t_n \xrightarrow{\tau} t, \text{ and for some } k, t_n \in C_n \forall n \geq k\}.$$

Kuratowski limits allow for a different characterization of Γ-limits.

Theorem 5. *Consider a sequence $F_n : T \to \mathbb{R}$. Let*

$$\operatorname{hypo}(F_n) = \{(t, x) \in T \times \mathbb{R} : F_n(t) \leq x\}.$$

Then $\Gamma\text{-}\lim F_n = F$ if and only if $\operatorname{Lihypo}(F_n) = \operatorname{Lshypo}(F_n) = \operatorname{hypo} F$.

Proof. See Dal Maso (1993), Theorem 4.16. □

In words, Γ-convergence of F_n to F is equivalent to the Kuratowski convergence of the hypo-graphs (the subset of $T \times \mathbb{R}$ that lies below the graph of F_n) of F_n to that of F. This gives an intuitive geometric characterization of Γ-convergence (and explains the equivalent name used in the literature of hypo-convergence/epi-convergence). Note that Kuratowski convergence is weaker than the notion of Hausdorff convergence, which is more familiar to economists.

Measure-theoretic preliminaries. Fix a measurable space $(\Omega, \mathcal{A})$ where $\mathcal{A}$ is a σ-algebra of subsets of Ω. As standard, call a map $\mu : \mathcal{A} \to [0, \infty)$ a *measure* if it is countably additive. A measurable function $f : \Omega \to \mathbb{R}$ is μ-integrable if

$$\sup\left\{ \int s \, d\mu : 0 \leq s \leq |f|, s : \Omega \to \mathbb{R} \text{ is simple} \right\} < \infty.$$

For $k \in \mathbb{N}$, call $\nu : \mathcal{A} \to \mathbb{R}^k$ a *vector measure* if for every sequence $(A_i)_{i=1}^{\infty}$ of pairwise disjoint sets it holds

$$\nu\left(\bigcup_{i=1}^{\infty} A_i \right) = \lim_{n\to\infty} \sum_{i=1}^{n} \nu(A_i),$$

where the limit on the right hand side is taken with respect to the norm defined by $\|\mathbf{x}\|_1 = \sum_{i=1}^{k} |x_i|$ for $\mathbf{x} = (x_i)_{i=1}^{k} \in \mathbb{R}^k$.

Given a vector measure v and $A \in \mathcal{A}$, let $|v|(A)$ be the measure given by

$$|v|(A) = \sup_{(B_i)_{i=1}^{m} \in \Pi(A)} \sum \|v(B_i)\|_1,$$

where $\Pi(A) = \{B = (B_i)_{i=1}^{\infty} : B \text{ is a partition of } A\}$. If $|v|(A) < \infty$ for every $A \in \mathcal{A}$, let $|v|$ denote the measure $|v| : \mathcal{A} \to \mathbb{R}$ defined by $A \mapsto |v|(A)$.

I provide some key facts about vector measures. Recall that measure v is non-atomic if for every A such that $v(A) > 0$ there exists $B \subset A$ such that $v(A) > v(B) > 0$. A vector measure v is non-atomic if $|v|$ is non-atomic. A collection $(A_\alpha)_{\alpha \in [0,1]}$ with $A_\alpha \in \mathcal{A}$ for every $\alpha \in [0,1]$ is a chain if $A_0 = \emptyset$, $A_1 = \Omega$ and $t \leq s \implies A_t \subseteq A_s$.

Lemma 2. *Let $v : \mathcal{A} \to \mathbb{R}^n$ be a vector measure.*

  (i) *$|v|$ is a measure;*
  (ii) *if v is a non-atomic vector measure, then the set*

  $$\{v(A) : A \in \mathcal{A}\},$$

  is compact and convex;
  (iii) *if v is non-atomic vector measure, then there exists a chain $(A_\alpha)_{\alpha \in [0,1]}$ such that*

  $$v(A_\alpha) = \alpha v(\Omega).$$

Proof. Point (i) is well-known so I omit the proof. See Diestel and Uhl (1977) for a proof. For point (ii), see Diestel and Uhl (1977) or Fryszkowski (2004). For a proof of (iii), see Fryszkowski (2004), Theorem 15. $\square$

The next result is well known, but I provide a proof for completeness.

Lemma 3. *Let μ be a real valued non-atomic measure. Then if $f : \Omega \to \mathbb{R}$ is integrable and $|f| \neq 0$ μ-a.s., the measure $\mu_{|f|}$ defined by*

$$\mu_{|f|}(A) = \int_A |f(\omega)| d\mu(\omega) \quad \forall A \in \mathcal{A},$$

is non-atomic.

Further, let μ be a measure and $f : \Omega \to \mathbb{R}$ μ-integrable. Define the measure μ_f by

$$\mu_f(A) = \int_A f(\omega) d\mu(\omega) \quad \forall A \in \mathcal{A}.$$

Then it holds that

$$|\mu_f|(A) = \int_A |f(\omega)| d\mu(\omega) = \mu_{|f|}(A),$$

for every $A \in \mathcal{A}$.

Proof. It is well known that $|f|$ is also integrable and that $|\mu_f|$ is a measure. Suppose that E is an atom for $\mu_{|f|}$, and consider the set $E' = \{|f| > 0\} \cap E$. Then E' is also an atom for $\mu_{|f|}$. Since $\mu_{|f|}$ is non-atomic, there exists $A \subset E'$ such that $0 < \mu_{|f|}(A) < \mu_{|f|}(E')$. Note that $f = 0$ μ-a.s. on either A or $E' \setminus A$ (otherwise $\mu_{|f|}(A) > 0$ and $\mu_{|f|}(E' \setminus A) > 0$, which contradicts the set E' being an atom). However, this contradicts the assumption that $|f|$ is positive on E'.

Now let $A \in \mathcal{A}$ and consider a partition $(B_i)_{i=1}^{\infty}$ of A. We have

$$\sum_{i=1}^{\infty} |\mu_f(B_i)| = \sum_{i=1}^{\infty} \left| \int_{B_i} f(\omega) d\mu(\omega) \right| \leq \sum_{i=1}^{\infty} \int_{B_i} |f(\omega)| d\mu(\omega)$$

$$= \int_A |f(\omega)| d\mu(\omega).$$

Thus $|\mu_f|(A) \leq \int_A |f(\omega)| d\mu(\omega)$.

Conversely, consider the partition of $A \in \mathcal{A}$ given by $B_1 = \{\omega \in A : f(\omega) \geq 0\}$ and $B_2 = \{\omega \in A : f(\omega) < 0\}$. Then by definition

$$\int |f(\omega)| d\mu(\omega) = \int_A f^+ d\mu(\omega) + \int_A f^- d\mu(\omega)$$

$$= |\mu_f(B_1)| + |\mu_f(B_2)| \leq |\mu_f|(A).$$

Hence $|\mu_f|(A) = \int_A |f(\omega)| d\mu(\omega)$ for every $A \in \mathcal{A}$ as desired. $\square$

Lemma 4. *Consider the vector measure $\lambda : \Sigma \to \mathbb{R}^2$ defined by*

$$A \mapsto \lambda(A) := \left(\int_A u(f) dP, P(A) \right).$$

Then λ is non-atomic.

Proof. Indeed, for every $A \in \Sigma$

$$|\lambda|(A) = \sup \sum_{i=1}^{\infty} P(A_i) + |P_{u \circ f}(A_i)|$$

$$= \sup \sum_{i=1}^{\infty} P(A_i) + \sum_{i=1}^{\infty} |P_{u \circ f}(A_i)|$$

$$= \sup \sum_{i=1}^{\infty} |P_{u \circ f}(A_i)| + P(A)$$

$$= |P_{u \circ f}|(A) + P(A),$$

where the supremum is over all Σ-measurable partitions $(A_i)_{i=1}^{\infty}$ of A. Thus $|\lambda| = |P_{u \circ f}| + P$. Now by Lemmas 3 and 3 $|P_{u \circ f}|$ is non-atomic (or identically zero, in which case the result is immediate), so that $|\lambda|$ is the sum of two non-atomic measures, hence it is non-atomic as well (e.g., see Johnson (1970), Theorem 1.2). $\square$

Proof of Theorem 1. Note that, since we are considering bounded acts, for every $f \in F$ there are x_f, y_f such that

$$y_f \leqslant f \leqslant x_f.$$

To prove the claim, there are three cases to consider. Define the following partition of $\mathbb{N}$:

$$N_1 := \left\{ n \in \mathbb{N} : \int u(f) dP = \int u(f) dP_n \right\},$$

$$N_2 := \left\{ n \in \mathbb{N} : \int u(f) dP < \int u(f) dP_n \right\},$$

and

$$N_3 := \left\{ n \in \mathbb{N} : \int u(f) dP > \int u(f) dP_n \right\}.$$

If $n \in N_1$ then we can just define $A_{n,f} \equiv S$ and $x_{n,f} \equiv x_f$. If $n \in N_2$ since

$$u(x_f) \geq \int u(f) dP_n,$$

we can apply Lemmas 4 and 2 to obtain a family $(A_\alpha)_{\alpha \in [0,1]}$ such that

$$\int u(f A_\alpha x_f) dP = \int_{A_\alpha} u(f) dP + (1 - P(A_\alpha)) u(x_f)$$

$$= \alpha \int u(f) dP + (1 - \alpha) u(x_f).$$

It follows that

$$\left\{ \int u(f A_\alpha x_f) dP : \alpha \in [0,1] \right\} = \left[\int u(f) dP, u(x_f) \right].$$

In particular, we have

$$\int u(f) dP_n \in \left[\int u(f) dP, u(x_f) \right],$$

so that $\int_{A_{\alpha_n}} u(f) dP + (1 - P(A_{\alpha_n})) u(x_f) = \int u(f) dP_n$ where

$$\alpha_n = 1 - \frac{\int u(f) dP - \int u(f) dP_n}{\int u(f) dP - u(x_f)}.$$

The last case to consider is $n \in N_3$, which can be dealt with symmetrically to the previous case (in particular, using y_f in place of x_f).

Hence for every n there must be α_n and $x \in \{x_f, y_f\}$ such that

$$\int_{A_{\alpha_n}} u(f)dP + (1 - P(A_{\alpha_n}))u(x) = \int u(f)dP_n.$$

In particular, the α_n's satisfy

$$\alpha_n = \begin{cases} 1 - \frac{\int u(f)dP - \int u(f)dP_n}{\int u(f)dP - u(y_f)} & \int u(f)dP - \int u(f)dP_n > 0, \\ 1 & \text{if } \int u(f)dP - \int u(f)dP_n = 0, \\ 1 - \frac{\int u(f)dP - \int u(f)dP_n}{\int u(f)dP - u(x_f)} & \int u(f)dP - \int u(f)dP_n(s) < 0. \end{cases}$$

Since $u(f) \in C_b(S, \mathbb{R})$ we have $\int u(f)dP_n \to \int u(f)dP$ it follows that $\alpha_n \to 1$. It follows that $P(A_{\alpha_n}) = \alpha_n \to 1$ as desired. $\quad\square$

Proof of Theorem 2. Given $(P_n)_{n=1}^{\infty}$ such that $P_n \overset{*}{\to} P$, for $n = 1, \ldots$ let $V_n : F \to \mathbb{R}$ be defined by

$$f \mapsto \int u(f)dP_n.$$

The next theorem is a type of Γ-convergence result for integral functionals.

Theorem 6. *For any* $P_n \overset{*}{\to} P$,

$$\Gamma\text{-}\lim V_n = V$$

Proof. See Theorem 14 in Lucchetti and Wets (1993). $\quad\square$

Proof of Theorem 2. Suppose that $\succeq$ has a preference for stability. By definition, this implies that for every $(F_n)_n$ there is $\varepsilon_n \to 0$ and a stable sequence $g_n = (f_n E_n x_n) \in \max_{u, P, \varepsilon_n}$.

In addition, it holds that

$$\int u(g_n)dP = \int u(f_n)dP_n.$$

Moreover, it must be that

$$\sup_{f \in F} \int u(f)dP_n = \sup_{g \in F_n} \int u(g)dP.$$

Indeed, by construction it holds

$$\left\{ \int u(f)dP_n : f \in F \right\} = \left\{ \int u(g)dP : g \in F_n \right\}.$$

Therefore, for every n it holds $f_n \in \max_{u, P_n, \varepsilon_n}(F)$. Since $(g_n)_n$ is stable, there is a subsequence $f_{n_k} \to f^* \in \max_{u, P}(F)$. I claim that this implies that

$$\lim_{n \to \infty} \sup_{f \in F} \int u(f)dP_n = \int u(f^*)dP. \tag{8}$$

First note that

$$\int u(f^*)dP \leq \liminf_{n \to \infty} \sup_{f \in F} \int u(f)dP_n.$$

This follows from Proposition 2.9 in Attouch (1984).

Finally, we also have that

$$\int u(f^*)dP \geq \limsup_{n \to \infty} \sup_{f \in F} \int u(f)dP_n.$$

Indeed,

$$\int u(f_n)dP_n \geq \sup_{f \in F} \int u(f)dP_n - \varepsilon_n, \tag{9}$$

which implies that

$$\limsup_{n \to \infty} \int u(f_n)dP_n \geq \limsup_{n \to \infty} \sup_{f \in F} \int u(f)dP_n,$$

but by definition this means that

$$\limsup_{n \to \infty} \int u(f_n)dP_n = \lim_{n \to \infty} \int u(f_{n_k})dP_{n_k},$$

for some subsequence n_k. By preference for stability, $f_{n_{k_j}} \to f' \in \max_{u, P}(F)$. But then since $\limsup_{n \to \infty} \int u(f_n)dP_n \geq \limsup_{n \to \infty} \sup_{f \in F} \int u(f)dP_n$ we get

$$\limsup_{j \to \infty} \int u(f_{n_{k_j}})dP_{n_{k_j}} \geq \limsup_{n \to \infty} \sup_{f \in F} \int u(f)dP_n.$$

By Γ-convergence,

$$\sup_{f \in F} \int u(f)dP = \int u(f')dP \geq \limsup_{j \to \infty} \int u(f_{n_{k_j}})dP_{n_{k_j}}$$

$$\geq \limsup_{n \to \infty} \sup_{f \in F} \int u(f)dP_n.$$

Hence (8) is proved.

Conversely, suppose that $\succeq$ is robust. Take $f^* \in \max_{u, P}(F)$. Since $\Gamma\text{-}\lim V_n = V$, there exists $f_n \to f^*$ such that $\lim_{n \to \infty} V_n(f_n) = V(f^*)$. Now I claim that for every $\varepsilon > 0$ there exists N_ε such that

$$f_n \in \max_{u, P_n, \varepsilon}(F) \text{ for all } n \geq N_\varepsilon.$$

To prove this claim, by contradiction assume that there is $\bar\varepsilon > 0$ and increasing map $\lambda : \mathbb{N} \to \mathbb{N}$ such that for every n

$$f_{\lambda(n)} \notin \max_{u, P_{\lambda(n)}, \bar\varepsilon}(F),$$

so that

$$\int u(f_{\lambda(n)})dP_{\lambda(n)} < \sup_{g \in F} \int u(g)dP_{\lambda(n)} - \bar\varepsilon,$$

but then we obtain

$$\int u(f^*)dP = \lim_{n \to \infty} \int u(f_{\lambda(n)})dP_{\lambda(n)} \leq \limsup_n \sup_{g \in F} \int u(g)dP_{\lambda(n)} - \bar\varepsilon$$

$$= \int u(f^*)dP - \bar\varepsilon,$$

a clear contradiction.

Now note that for every $\varepsilon > 0$ we can modify the sequence $(f_n)_n$ into the sequence $(f_n^\varepsilon)_n$ so that for every n it holds

$$f_n^\varepsilon \in \max_{u, P, \varepsilon}(F).$$

But then defining the double-indexed sequence $g_{k,n}$ by

$$g_{k,n} \equiv \left(\int u(f_n^{\frac{1}{k}})dP_n, f_n^{\frac{1}{k}}, \int u(f_n^{\frac{1}{k}})dP \right),$$

which satisfies for every k

$$g_{k,n} \to \left(\int u(f^*)dP, f^*, \int u(f^*)dP \right),$$

where convergence is in the topological space $\mathbb{R} \times C_b(S, X) \times \mathbb{R}$ endowed with the product topology. Hence by Lemma 1 there exists $\iota : \mathbb{N} \to \mathbb{N}$ increasing and with $\lim_{n \to \infty} \iota(n) = \infty$ such that $\lim_{n \to \infty} g_{\iota(n),n} = \left(\int u(f^*)dP, f^*, \int u(f^*)dP \right)$. Therefore,

$$f_n^{\frac{1}{\iota(n)}} \in \max_{u, P_n, \frac{1}{\iota(n)}}(F), \frac{1}{\iota(n)} \to 0.$$

Set $f_n^{\varepsilon_n} \equiv f_n^{\frac{1}{\iota(n)}}$ and note that $f_n^{\varepsilon_n} \in \max_{u, P_n, \varepsilon_n}$. Note that $f_n^{\varepsilon_n} \to f^*$ and $\int u(f_n^{\varepsilon_n})dP_n \to \int u(f^*)dP$.

I now claim that there exists a sequence $g_n \in F_n$ such that $g_n = f_n^{\varepsilon_n} E_n x_n$, $P(E_n) \to 1$, x_n eventually takes only two different values and

$$\int u(f_n^{\varepsilon_n} E_n x_n)dP = \int u(f_n^{\varepsilon_n})dP_n.$$

Indeed, since $\int u(f_n^{\varepsilon_n})dP_n \to \int u(f^*)dP$, there are x^*, y^* and N such that for every $n \geq N$

$$u(y^*) \leq \int u(f_n^{\varepsilon_n})dP_n \leq u(x^*).$$

To see this, let x^*, y^* be constant acts such that $x^* \geqslant f^* \geqslant y^*$. Note that we have

$$\int u(f^*)dP \in (u(y^*), u(x^*)),$$

so that by definition of convergence there must be N such that for every $n \geq N$

$$u(y^*) \leq \int u(f_n^{\varepsilon_n})dP_n \leq u(x^*),$$

as desired. Given this result, I proceed as in the proof of Theorem 1. Define the vector measure λ^n by

$$A \mapsto (\int_A u(f_n^{\varepsilon_n})dP, P(A)) \quad \text{for every } A \in \Sigma.$$

By the same reasoning as in the Proof of Theorem 1, λ^n is non-atomic. Hence for every n there is a chain $(E_\alpha^n)_{\alpha \in [0,1]}$ such that for $x \in \{x^*, y^*\}$ it holds

$$\int_{E_\alpha^n} u(f_n^{\varepsilon_n})dP + (1 - P(E_\alpha^n))u(x) = \alpha \int u(f_n^{\varepsilon_n})dP + (1 - \alpha)u(x).$$

Thus we can find sequences $(E_n)_n$ and $(x_n)_n$ with $x_n \in \{x^*, y^*\}$ such that $\int u(f_n^{\varepsilon_n} E_n x_n)dP = \int u(f_n^{\varepsilon_n})dP_n$ and

$$1 - P(A_n) = \begin{cases} \frac{\int u(f_n^{\varepsilon_n})dP - \int u(f_n^{\varepsilon_n})dP_n}{\int u(f_n^{\varepsilon_n})dP - u(y^*)} & \int u(f_n^{\varepsilon_n})dP - \int u(f_n^{\varepsilon_n})dP_n > 0, \\ 0 & \text{if } \int u(f_n^{\varepsilon_n})dP - \int u(f_n^{\varepsilon_n})dP_n = 0, \\ \frac{\int u(f_n^{\varepsilon_n})dP - \int u(f_n^{\varepsilon_n})dP_n}{\int u(f_n^{\varepsilon_n})dP - u(x^*)} & \int u(f_n^{\varepsilon_n})dP - \int u(f_n^{\varepsilon_n})dP_n < 0. \end{cases}$$

But note that $\int u(f_n^{\varepsilon_n})dP - \int u(f_n^{\varepsilon_n})dP_n \to 0$. Indeed, we know that $\int u(f_n^{\varepsilon_n})dP_n \to \int u(f^*)dP$. Moreover, $\int u(f_n^{\varepsilon_n})dP \to \int u(f^*)dP$. Hence

$$1 - P(A_n) \to 0,$$

as desired.

Thus letting $g_n \equiv f_n^{\varepsilon_n} E_n x_n$ we have $g_n \in \max_{u,P,\varepsilon_n}(F_n)$. Finally, it holds that for a subsequence $f_{n_k}^{\varepsilon_{n_k}}$

$$\int \|f^* - f_{n_k}^{\varepsilon_{n_k}} E_{n_k} x_{n_k}\|dP \to 0.$$

Indeed,

$$\int \|f^* - f_n^{\varepsilon_n} E_n x_n\|dP \leq \int \|f_n - f_n E_n x_n\|dP + \int \|f_n - f^*\|dP. \quad (10)$$

Since $f_n^{\varepsilon_n} \to f^*$ uniformly, it follows that

$$\int \|f_n^{\varepsilon_n} - f^*\|dP \to 0.$$

Finally, note that we have

$$\int \|f_n^{\varepsilon_n} - f_n^{\varepsilon_n} E_n x_n\|dP = \int \|\mathbf{0} E_n x_n\|dP = \int \|x_n\| 1_{E_n^c} dP,$$

where $\mathbf{0}$ denotes the zero vector. If necessary, by passing to a subsequence, $1_{A_n^c} \to \mathbf{0}$ P-a.s. thus by applying Lebesgue's dominated convergence theorem we get

$$\int \|f_{n_k}^{\varepsilon_{n_k}} - f_{n_k}^{\varepsilon_{n_k}} E_{n_k} x_{n_k}\|dP \to 0,$$

which by (10) gives

$$\int \|f^* - f_{n_k}^{\varepsilon_{n_k}} E_{n_k} x_{n_k}\|dP \to 0,$$

as desired. $\quad\square$

Proof of Proposition 1. Since F is a compact subset of $C_b(S, X)$ the sequence V_n is equi-coercive (cf. Dal Maso (1993), Definition 1.12). Now since V_n Γ-converges to V, by applying Theorem 7.8 in Dal Maso (1993) the result follows. If the optimum is unique, then it suffices to apply Corollary 7.24 in Dal Maso (1993). $\quad\square$

Proof of Corollary 1. The "if" part is immediate.

Conversely, if f_n is a sequence of ε_n acts that converges to an optimal act f^*, then by applying Corollary 7.20 in Dal Maso (1993) we get that

$$\int u(f_n)dP_n \to \int u(f^*)dP.$$

By using the same reasoning as in the proof of Theorem 2, one can construct a sequence $g_n = f_n E_n x_n$ such that $g_n \in \max_{u,P,\varepsilon_n}(F_n)$ and for a subsequence g_{n_k} satisfies

$$\int \|g_{n_k} - f^*\|dP \to 0,$$

as desired. $\quad\square$

Proof of Theorem 3.

Lemma 5. $V : F \to \mathbb{R}$ *is continuous.*

Proof. Omitted. $\quad\square$

Proof of Theorem 3. The proof of this result is based on standard techniques; see for example Bonnans and Shapiro (2013), Proposition 4.12 or Battauz et al. (2015). Consider any $h_n \downarrow 0$. Note that by definition it holds $\int u(f^*)d((1 - h_n)P + h_n Q) \leq W((1 - h_n)P + Q h_n)$ and $\int u(f^*)d((1 - h_n)P + h_n Q)$. Moreover, letting $\varepsilon_n = \frac{h_n}{n}$, any sequence $(f_{\varepsilon_n})_n$ satisfies

$$W((1 - h_n)P + Q h_n) \leq \int u(f_{\varepsilon_n})d((1 - h_n)P + Q h_n) + \varepsilon_n.$$

Also note that

$$\frac{\int u(f_{\varepsilon_n})d((1 - h_n)P + Q h_n) - \int u(f_{\varepsilon_n})dP}{h_n} = \int u(f_{\varepsilon_n}^*)d(Q - P).$$

Thus we get the following inequalities

$$\frac{\int u(f^*)d((1 - h_n)P + h_n Q) - \int u(f^*)dP}{h_n} - \int u(f^*)d(Q - P)$$

$$\leq \frac{W((1 - h_n)P + Q h_n) - W(P)}{h_n} - \int u(f^*)d(Q - P)$$

$$\leq \int u(f_{\varepsilon_n}^*)d(Q - P) + \frac{\varepsilon_n}{h_n} - \int u(f^*)d(Q - P).$$

Then since we have robustness it follows that $f_{\varepsilon_{h_n}} \to f^*$. Indeed, because the optimal act is unique, by Theorem 2 and Corollary 7.17 in Dal Maso (1993), any convergent subsequence $(f_{\varepsilon_{h_n}})_k$ converges to f^*, which implies that $f_{\varepsilon_{h_n}} \to f^*$. Hence we find that

$$\int u(f_{\varepsilon_{h_n}}^*)d(Q - P) + \frac{\varepsilon_{h_n}}{h_n} - \int u(f^*)d(Q - P) \to 0.$$

Indeed, by Lemma 5

$$\int u(f_{\varepsilon_{h_n}}^*)d(Q - P) \to \int u(f^*)d(Q - P),$$

so that

$$\int u(f_{\varepsilon_{h_n}}^*)d(Q - P) + \frac{\varepsilon_{h_n}}{h_n} - \int u(f^*)d(Q - P) \to 0.$$

Since this holds for any $h_n \downarrow 0$, we can conclude that

$$\lim_{h \downarrow 0} \frac{W((1 - h)P + Qh) - W(P)}{h} = \int u(f^*)dQ - \int u(f^*)dP. \quad\square$$

Proof of Proposition 3. First, since C is weak*-compact, there are $Q_1, Q_2 \in C$ such that

$$\sup_{Q \in C} \int u(f_1)dQ = \int u(f_1)dQ_1,$$

and $\sup_{Q \in C} \int u(f_2)dQ = \int u(f_2)dQ_2$. It follows that for $i = 1, 2$ there exist $f_i E_i x_i \in F_{Q_i}$ such that $f_i E_i x_i \succeq_i g$, for all $g \in F_C^i$. Because

$\int u(f_1)dP_1 = \int u(f_2)dP_2$, $m(u,P_1,F) \geq m(u,P_2,F) \iff \int u(f_1)dQ_1 \geq \int u(f_2)dQ_2$. Moreover, by assumption $\succcurlyeq_1$ and $\succcurlyeq_2$ admit certainty equivalents. Therefore, there exist $x,y \in X$ such that $f_1 E_1 x_1 \sim_1 x$ and $f_2 E_2 x_2 \sim_2 y$, so that

$$\int u(f_1)dQ_1 = \int u(f_1 E_1 x_1)dP_1 = u(x),$$

and

$$\int u(f_2)dQ_2 = \int u(f_2 E_2 x_2)dP_2 = u(y),$$

from which the result follows.

Proof of Proposition 4. Observe that

$$-\sup_{Q \in C}\int u(f^*)dQ = \inf_{Q \in C}\int -u(f^*)dQ.$$

It follows that one can use Theorem 1 in Luenberger (1997) (p. 217). Indeed, note that all the assumptions of this result are satisfied: the map $Q \mapsto \int -u(f^*)dQ$ is affine and thus convex, C is a convex subset of $ca(\Sigma)$ (by well known results), and $R(P\|P) = 0 < K$. Hence there exists θ decreasing with K such that

$$\inf_{Q \in C}\int -u(f^*)dQ = \inf_{Q \in \Delta}\int -u(f^*)dQ + \theta R(Q\|P), \text{ for some } \theta \geq 0.$$

Since the map $s \mapsto -u(f^*(s))$ is a bounded measurable function, by applying a well-known variational formula (e.g., see Dupuis and Ellis (1997), Proposition 1.4.2, p. 27) then we obtain,

$$\inf_{Q \in \Delta}\int -u(f^*)dQ + \theta R(Q\|P) = -\theta \log\left(\int e^{\frac{1}{\theta}u(f^*)}dP\right),$$

so that

$$\sup_{Q \in C}\int u(f^*)dQ - \int u(f^*)dP = \theta\log\left(\int e^{\frac{1}{\theta}u(f^*)}dP\right) - \int u(f^*)dP.$$

as desired. ☐

Proof of Proposition 5. The optimization problem for every $P \in C$ with parameters (μ,σ^2) can be written as

$$V(\mu,\sigma^2) = \max_{a \in [0,\infty)} -e^{\frac{\sigma^2}{2}-a\mu} - e^{a-\bar{c}}.$$

The solution is given by the first order condition:

$$\mu e^{\frac{\sigma^2}{2}-a\mu} = e^{a-\bar{c}}.$$

We obtain the unique solution for P with parameters (μ,σ^2)

$$a_P = \frac{\ln(\mu) + \frac{\sigma^2}{2} + \bar{c}}{\mu + 1}.$$

So, we have

$$V(\mu,\sigma^2) = -e^{\frac{\sigma^2}{2}-a_P\mu} - e^{a_P-\bar{c}}.$$

Since the objective function is strictly increasing in μ and strictly decreasing in σ^2, it is straightforward to see that

$$\frac{\partial V(\mu,\sigma^2)}{\partial\mu} > 0, \tag{11}$$

and

$$\frac{\partial V(\mu,\sigma^2)}{\partial\sigma^2} < 0. \tag{12}$$

It follows that $\sigma_2^2 > \sigma_1^2$ and $V(P_1) = V(P_2)$ imply that $\mu_2 > \mu_1$. Moreover, together (11) and (12) also imply that

$$\sup_{(\mu,\sigma^2)}\left\{-e^{\frac{\sigma^2}{2}-a_P\mu} - e^{a_P-\bar{c}}\right\} - V(\mu,\sigma^2) = e^{\frac{\sigma^2}{2}-a_P\mu} - e^{\frac{\sigma^2}{2}-a_P\bar{\mu}}.$$

Hence we obtain

$$m(v,P_1,A) \geq m(v,P_2,A)$$
$$\iff e^{\frac{\sigma_1^2}{2}-a_{P_1}\mu_1} - e^{\frac{\sigma^2}{2}-a_{P_1}\bar{\mu}} \geq e^{\frac{\sigma_2^2}{2}-a_{P_2}\mu_2} - e^{\frac{\sigma^2}{2}-a_{P_2}\bar{\mu}},$$

which is equivalent to

$$e^{\frac{\sigma_1^2}{2}-a_{P_1}\mu_1} - e^{\frac{\sigma_2^2}{2}-a_{P_2}\mu_2} \geq e^{\frac{\sigma^2}{2}-a_{P_1}\bar{\mu}} - e^{\frac{\sigma^2}{2}-a_{P_2}\bar{\mu}},$$

which in turn is implied by convexity of e^x combined with (7). We can therefore conclude that $m(v,P_1,A) \geq m(v,P_2,A)$ as desired. ☐

Proof of Proposition 6. It is enough to check that for $i = 1,2$ and every $Q \in C$

$$\limsup_{h\downarrow 0}\sup_{\alpha\in[0,1]}\frac{\mathbb{E}_{hQ+(1-h)P_i}u(\alpha\omega+(1-\alpha)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i))}{h}$$
$$= E_Q u(\alpha_i\omega+(1-\alpha_i)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i)).$$

Indeed, this would imply that

$$m(u,P_i,F) = \sup_{Q\in C}\mathbb{E}_Q u(\alpha\omega+(1-\alpha)) - \mathbb{E}_{P_i}u(\alpha\omega+(1-\alpha)),$$

so that by applying the same reasoning as in Proposition we get that there exists $\theta \geq 0$ decreasing with the size K of the neighborhood C such that

$$\sup_Q \mathbb{E}_{Q\in C}u(\alpha_i\omega+(1-\alpha_i)) = \theta\log\left(\mathbb{E}_P e^{\frac{1}{\theta}u(\alpha_i\omega+(1-\alpha_i))}\right),$$

from which the desired result follows.

To prove the claim, note that we have the following inequality for any $h_n \downarrow 0$

$$\frac{\mathbb{E}_{h_nQ+(1-h_n)P_i}u(\alpha_i\omega+(1-\alpha_i)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i))}{h_n}$$
$$\leq \sup_{\alpha\in[0,1]}\frac{\mathbb{E}_{h_nQ+(1-h_n)P_i}u(\alpha\omega+(1-\alpha)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i))}{h_n},$$

so that,

$$\mathbb{E}_Q u(\alpha_i\omega+(1-\alpha_i)) - \mathbb{E}_{P_i}u(\alpha_i\omega+(1-\alpha_i))$$
$$= \lim_{n\to\infty}\frac{\mathbb{E}_{h_nQ+(1-h_n)P_i}u(\alpha_i\omega+(1-\alpha_i)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i))}{h_n}$$
$$\leq \sup_{\alpha\in[0,1]}\frac{\mathbb{E}_{h_nQ+(1-h_n)P_i}u(\alpha\omega+(1-\alpha)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i))}{h_n}.$$

Moreover, for any sequence $(\alpha_i^n)_n$ of wealth allocations optimal for $h_nQ+(1-h_n)P_i$ (they exist by compactness and continuity), because α_i is the unique optimal allocation for the belief α, by the same reasoning as in Proposition 1 it holds that $\alpha_i^n \to \alpha_i$ which implies

$$\mathbb{E}_Q u(\alpha_i\omega+(1-\alpha_i)) - \mathbb{E}_{P_i}u(\alpha_i\omega+(1-\alpha_i))$$
$$= \lim_{n\to\infty}\frac{\mathbb{E}_{h_nQ+(1-h_n)P_i}u(\alpha_i^n\omega+(1-\alpha_i^n)) - E_{P_i}u(\alpha_i^n\omega+(1-\alpha_i^n))}{h_n}$$
$$\geq \sup_{\alpha\in[0,1]}\frac{\mathbb{E}_{h_nQ+(1-h_n)P_i}u(\alpha\omega+(1-\alpha)) - E_{P_i}u(\alpha_i\omega+(1-\alpha_i))}{h_n}.$$

so that

$$m(u,P_i,F) = \sup_{Q\in C}\mathbb{E}_Q u(\alpha\omega+(1-\alpha)) - \mathbb{E}_{P_i}u(\alpha\omega+(1-\alpha)),$$

as desired. ☐

References

Ackerman, F., Stanton, E.A., Bueno, R., 2010. Fat tails, exponents, extreme uncertainty: Simulating catastrophe in DICE. Ecol. Econom. 69 (8), 1657–1665.

Aliprantis, C., Border, K., 2006. Infinite Dimensional Analysis. Springer.

Attouch, H., 1984. Variational Convergence for Functions and Operators, Vol. 1. Pitman Advanced Publishing Program.

Battauz, A., De Donno, M., Ortu, F., 2015. Envelope theorems in Banach lattices and asset pricing. Math. Financ. Econ. 9 (4), 303–323.

Berger, J.O., Insua, D.R., Ruggeri, F., 2000. Robust Bayesian Analysis. Springer.

Berger, L., Marinacci, M., 2020. Model uncertainty in climate change economics: A review and proposed framework for future research. Environ. Resour. Econ. 77 (3), 475–501.

Bommier, A., Fabre, A., Goussebaïle, A., Heyen, D., 2021. Disagreement aversion. Available at SSRN, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3964182.

Bonnans, J.F., Shapiro, A., 2013. Perturbation Analysis of Optimization Problems. Springer Science & Business Media.

Braides, A., 2002. Gamma-Convergence for Beginners, Vol. 22. Clarendon Press.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L., Stanca, L., 2019. Affine Gateaux Differentials and the Von Mises Statistical Calculus. Mimeo.

Dal Maso, G., 1993. An Introduction To Γ-Convergence, Vol. 8. Springer Science & Business Media.

De Finetti, B., 1961. The Bayesian approach to the rejection of outliers. In: Proceedings of the Fourth Berkeley Symposium on Probability and Statistics, Vol. 1. University of California Press Berkeley, pp. 199–210.

Dekel, E., Fudenberg, D., Morris, S., 2006. Topologies on types. Theor. Econ. 1 (3), 275–309.

Dempster, A.P., 1975. A subjectivist look at robustness. Bull. Int. Stat. Inst. 46, 349–374.

Diestel, J., Uhl, J., 1977. Vector Measures. American Mathematical Society.

Dunford, N., Schwartz, J.T., 1958. Linear Operators, Part 1: General Theory, Vol. 7. John Wiley & Sons.

Dupuis, P., Ellis, R.S., 1997. A Weak Convergence Approach To the Theory of Large Deviations, Vol. 313. John Wiley & Sons.

Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. Q. J. Econ. 75 (4), 643–669.

Fishburn, P.C., Murphy, A.H., Isaacs, H.H., 1968. Sensitivity of decisions to probability estimation errors: A reexamination. Oper. Res. 16 (2), 254–267.

Fryszkowski, A., 2004. Fixed Point Theory for Decomposable Sets, Vol. 2. Springer Science & Business Media.

Gilboa, I., Marinacci, M., 2016. Ambiguity and the Bayesian paradigm. In: Readings in Formal Epistemology. Springer, chapter 21.

Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. J. Math. Econ. 18 (2), 141–153.

Gollier, C., Jullien, B., Treich, N., 2000. Scientific progress and irreversibility: an economic interpretation of the 'Precautionary Principle'. J. Public Econom. 75 (2), 229–253.

Good, I., 1971. The probabilistic explication of information, evidence, surprise, causality, explanation, and utility. In: Foundations of Statistical Inference, Vol. 57. Hold, Rienhart and Winston, Toronto, pp. 108–127.

Hampel, F.R., 1971. A general qualitative definition of robustness. Ann. Math. Stat. 42 (6), 1887–1896.

Hampel, F.R., 1974. The influence curve and its role in robust estimation. J. Am. Stat. Assoc. 69 (346), 383–393.

Hansen, L., Sargent, T.J., 2001. Robust control and model uncertainty. Amer. Econ. Rev. 91 (2), 60–66.

Huber, P.J., et al., 1964. Robust estimation of a location parameter. Ann. Math. Stat. 35 (1), 73–101.

Hyndman, K., Ozbay, E.Y., Schotter, A., Ehrblatt, W.Z., 2012. Convergence: an experimental study of teaching and learning in repeated games. J. Eur. Econom. Assoc. 10 (3), 573–604.

Ibragimov, M., Ibragimov, R., Walden, J., 2015. Heavy-Tailed Distributions and Robustness in Economics and Finance, Vol. 214. Springer.

Ikefuji, M., Laeven, R.J., Magnus, J.R., Muris, C., 2020. Expected utility and catastrophic risk in a stochastic economy–climate model. J. Econometrics 214 (1), 110–129.

Johnson, R.A., 1970. Atomic and nonatomic measures. Proc. Amer. Math. Soc. 25 (3), 650–655.

Kadane, J.B., Chuang, D.T., 1978. Stable decision problems. Ann. Statist. 6 (5), 1095–1110.

Kopylov, I., 2010. Simple axioms for countably additive subjective probability. J. Math. Econom. 46 (5), 867–876.

Lucchetti, R., Wets, R.J., 1993. Convergence of minima of integral functionals, with applications to optimal control and stochastic optimization. Stat. Risk Model. 11 (1), 69–84.

Luenberger, D.G., 1997. Optimization By Vector Space Methods. John Wiley & Sons.

Marinacci, M., 2015. Model uncertainty. J. Eur. Econom. Assoc. 13 (6), 1022–1100.

Meinhold, R.J., Singpurwalla, N.D., 1989. Robustification of Kalman filter models. J. Amer. Statist. Assoc. 84 (406), 479–486.

Meinshausen, M., Meinshausen, N., Hare, W., Raper, S.C., Frieler, K., Knutti, R., Frame, D.J., Allen, M.R., 2009. Greenhouse-gas emission targets for limiting global warming to 2 degrees C. Nature 458 (7242), 1158–1162.

Müller, U.A., Dacorogna, M.M., Pictet, O.V., 1998. Heavy tails in high-frequency financial data. In: A Practical Guide To Heavy Tails: Statistical Techniques and Applications. Birkhäuser Boston, pp. 55–78, chapter 3.

Murphy, K.P., 2012. Machine Learning: A Probabilistic Perspective. MIT Press.

Ok, E.A., 2011. Real analysis with economic applications. In: Real Analysis with Economic Applications. Princeton University Press.

Pierce, D.A., Folks, J.L., 1969. Sensitivity of Bayes procedures to the prior distribution. Oper. Res. 17 (2), 344–350.

Rubinstein, A., 1989. The electronic mail game: Strategic behavior under" almost common knowledge". Am. Econ. Rev. 79 (3), 385–391.

Salinetti, G., 1994. Stability of Bayesian decisions. J. Stat. Plan. Inference 40 (2–3), 313–329.

Savage, L., 1954. Foundations of Statistics. Wiley.

Savage, L.J., Edwards, W., Lindman, H., 1963. Bayesian statistical inference for psychological research. Psychol. Rev. 70 (3), 193.

Schmeidler, D., 1989. Subjective probability and expected utility without additivity. Econometrica 57 (3), 571–587.

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

Taleb, N.N., Cirillo, P., 2019. Branching epistemic uncertainty and thickness of tails. Available at arXiv https://arxiv.org/abs/1912.00277.

Weitzman, M.L., 2007. Subjective expectations and asset-return puzzles. Amer. Econ. Rev. 97 (4), 1102–1130.

Weitzman, M.L., 2011. Fat-tailed uncertainty in the economics of catastrophic climate change. Rev. Environ. Econ. Policy 5 (2), 275–292.