MAXMIN UTILITY OR MINMAX REGRET? CONSISTENCY AND FALSIFIABILITY UNDER MULTIPLE PRIORS

Andriy Zapechelnyuk

Abstract. A criterion of choice under ambiguity is consistent if its choices do not depend on which utility function is used to represent the decision maker's preferences in the absence of ambiguity. The criterion is falsifiable if observed choices can, at least sometimes, refute a hypothesis about these preferences. Both matter in empirical and experimental work that estimates utility functions. We characterize consistent criteria and give a simple sufficient condition for unfalsifiability. In particular, minmax regret is consistent and falsifiable, while maxmin utility is neither. The problem with maxmin utility disappears if utility is normalized or measured relative to a reference point.

JEL Classification: D81, C18

Keywords: Ambiguity, multiple priors, expected utility, maxmin utility, minmax regret, consistency, falsifiability

November 23, 2026

1. Introduction

The purpose of this note is to clarify that for some criteria of choice under ambiguity, primitive utility functions that represent the same expected-utility preferences need not lead to the same optimal choices under ambiguity. In the extreme case, observed choices under ambiguity may reveal nothing at all about the underlying preferences.

Consider a decision maker who is endowed with a set of priors over states and chooses among several actions. The decision maker's preferences over choices with a single prior (the standard expected utility setting) are captured by a utility function. Two

Zapechelnyuk: School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK. *E-mail:* andy@zapechelnyuk.com.

The author is grateful for helpful comments to Benson Leung and Karl Schlag.

utility functions represent the same preferences under expected utility if they do not change the comparison between choices, that is, for every prior, the difference in the expected utility between every pair of actions is preserved, up to rescaling.

We study a general class of criteria of choice under ambiguity, as in Cerreia-Vioglio et al. (2011). This class includes most ambiguity averse preferences that appear in the literature. We consider two basic properties of a criterion of choice under ambiguity. The first property is *consistency*, saying that the comparison between choices under ambiguity is preserved, up to rescaling, for all primitive utility functions that represent the same preferences under expected utility. A second, much weaker property is *falsifiability*, saying that the decision maker's observed choices should, at least sometimes, reveal something about her preferences. Namely, for some preferences and some set of priors, there exists an undominated choice that is not optimal under any utility function that represents these preferences, and thus, observing this choice falsifies the hypothesis that the decision maker has these preferences. So, unfalsifiability is an extreme form of the lack of consistency, where not only the comparison of choices is not preserved, but also every undominated choice can be made optimal under some utility function representing the given preferences.

We obtain two results. First, we characterize consistent criteria. Briefly, a criterion is consistent if changing the utility representation transforms the evaluations of all choices affinely under each single prior, and it does not change which prior is the worse one among each pair of priors. Second, we provide a simple sufficient condition for unfalsifiability. We apply these results to a range of criteria. In particular, it turns out that minmax regret (Savage, 1951) and reference-based maxmin utility (Bewley, 1986) are consistent and falsifiable, competitive ratio (Sleator and Tarjan, 1985) is falsifiable but not consistent, and standard maxmin utility (Gilboa and Schmeidler, 1989) is neither consistent nor falsifiable. Intuitively, the reason why maxmin utility fails both criteria is that the comparison under maxmin utility has an implicit absolute reference point of zero. When a preference-preserving transformation is applied to a utility function, this reference point does not change, but the comparison of the choices to it does. Rather than arguing that maxmin utility is not a suitable criterion, we point out that one must be aware of the highlighted issues and either accept that the absolute comparison to zero is valid and meaningful in their framework, or change

the criterion, say, to its reference-based or normalized counterpart to avoid the issue altogether.

Related Literature. This paper contributes to the literature on choice under ambiguity. A variety of criteria of choice under ambiguity are proposed in this literature, including maxmin utility (Gilboa and Schmeidler, 1989), minmax regret (Savage, 1951; Milnor, 1954; Stoye, 2011), variational preferences (Maccheroni et al., 2006), multiplier preferences (Hansen and Sargent, 2001), smooth ambiguity preferences (Klibanoff et al., 2005), maxmin safety (Gulko and Leung 2013, also known as maxmin joy, Hayashi 2008), reference-based maxmin utility (Bewley, 1986; Ortoleva, 2010), and the competitive ratio (Sleator and Tarjan, 1985). These criteria belong to the class of uncertainty averse preferences of Cerreia-Vioglio et al. (2011), and most of them have axiomatic foundations. In this paper, we do not introduce a new criterion, nor do we provide a new axiomatic foundation for an existing one. Instead, we point out two properties, consistency and falsifiability, that matter when these criteria are applied in empirical and experimental studies that elicit preferences and estimate a utility function.

Viewing choice under ambiguity as a game against nature, consistency requires the decision maker's choices to depend on her payoffs only up to strategic equivalence (e.g., Moulin and Vial, 1978), that is, up to a positive rescaling and the addition of a constant that depends on nature's choice of state. Under ignorance (the set of degenerate priors), this is the content of the linearity and column linearity axioms of Milnor (1954) (see also Luce and Raiffa, 1957, Chapter 13). We extend this requirement to arbitrary sets of priors and to the general class of criteria of Cerreia-Vioglio et al. (2011), and motivate it by identification: state-specific constants are not pinned down by preferences under expected utility, so choices under ambiguity should not depend on them either. This issue does not arise in the framework of Anscombe and Aumann (1963), which underlies most axiomatizations of ambiguity averse preferences. This is because in Anscombe and Aumann (1963), utility does not depend on the state directly (only indirectly, via probabilities of consequences) and is unique up to a common positive affine transformation.

Falsifiability is related to the revealed preference literature that studies testable implications of models of choice under risk and uncertainty (see Echenique, 2020, for a survey). For example, Chambers et al. (2016) test whether the data on asset demand

are consistent with maxmin expected utility and variational preferences. We ask a more general question: given a specified criterion of choice, the decision maker's preferences under expected utility, and a set of priors, is there any undominated choice that the criterion rules out? Proposition 2 shows that, in particular, for maxmin utility and variational preferences, there is none. This is reminiscent of the classical observation that, when utility depends on the state, beliefs and utilities cannot be separately identified from choices (Karni et al., 1983; Karni, 1985). In our setting, the freedom to add state-specific constants to the utility function is enough to rationalize every undominated choice when the criterion is unfalsifiable.

2. Model

A decision maker (DM) chooses an action x from a finite set X. A state of nature ω is drawn from a finite set of states Ω. The DM has preferences over actions and states, which are represented by a utility function $u : X \times \Omega \to \mathbb{R}$. Let $\bar{\mathcal{U}}$ be the set of all utility functions. We linearly extend the definition of $u \in \bar{\mathcal{U}}$ to the mixed domain $\Delta(X) \times \Delta(\Omega)$,

$$u(\sigma, p) = \sum_{x \in X} \sum_{\omega \in \Omega} u(x, \omega)\sigma(x)p(\omega) \quad \text{for all } (\sigma, p) \in \Delta(X) \times \Delta(\Omega),$$

and identify each action $x \in X$ with the degenerate mixed action δ_x, and each state $\omega \in \Omega$ with the degenerate prior δ_ω.

The DM is endowed with a set of priors $P \subseteq \Delta(\Omega)$ and is ambiguity averse. Let $\mathcal{P}$ be the set of all nonempty compact subsets of $\Delta(\Omega)$. Following Cerreia-Vioglio et al. (2011), we consider a general class of criteria that capture the DM's preferences under ambiguity, described by an *ambiguity operator* $G : \mathbb{R} \times \Delta(\Omega) \times \bar{\mathcal{U}} \to \mathbb{R}$.[1] Assume that $G(t, p, u)$ is continuous in (t, p), strictly increasing in t, and $G(\mathbb{R}, p, u) = \mathbb{R}$ for all $(p, u) \in \Delta(\Omega) \times \bar{\mathcal{U}}$. Then, given a utility function u, an ambiguity operator G, and a set of priors P, the DM's problem is

$$\max_{\sigma \in \Delta(X)} \min_{p \in P} G(u(\sigma, p), p, u). \tag{1}$$

[1]The competitive ratio criterion of Sleator and Tarjan (1985) requires utility functions to satisfy $\max_{x \in X} u(x, p) > 0$ for every $p \in \Delta(\Omega)$, so in this case we consider this restricted domain of utility functions. As a result, Proposition 2 does not apply to this case as stated. Instead, we deal with this case directly in Appendix A.3.

A defining feature of criterion (1) is that a mixed action σ only enters in the criterion through the utility $u(\sigma, p)$. Representation (1) has the form of the uncertainty averse preferences of Cerreia-Vioglio et al. (2011) and encompasses most ambiguity averse preferences in the literature, as listed in Table 1, where we use notation

$$\underline{u}(p) = \min_{x \in X} u(x, p) \quad \text{and} \quad \bar{u}(p) = \max_{x \in X} u(x, p).$$

TABLE 1. Ambiguity operators.

Criterion	$G(t, p, u)$	Notes	Reference
Maxmin utility	t		Gilboa and Schmeidler (1989)
Minmax regret	$t - \bar{u}(p)$		Savage (1951); Stoye (2011)
Maxmin safety	$t - \underline{u}(p)$		Hayashi (2008); Gulko and Leung (2013)
Competitive ratio	$\dfrac{t}{\bar{u}(p)}$	$\bar{u}(p) > 0$ for all p	Sleator and Tarjan (1985)
Normalized maxmin utility	$\dfrac{t - \underline{u}(p)}{\bar{u}(p) - \underline{u}(p)}$	$\bar{u}(p) > \underline{u}(p)$ for all p	Maxmin utility with normalization of Dhillon and Mertens (1999)
Reference-based utility	$t - u(x_0, p)$	$x_0 \in X$ is a reference action	Bewley (1986); Ortoleva (2010)
Multiplier preferences	$t + \theta R(p\|\bar{p})$	$\theta > 0$, R is relative entropy, $\bar{p}$ has full support	Hansen and Sargent (2001)
Variational preferences	$t + \phi(p)$	$\phi \geq 0$ continuous, $\min_p \phi(p) = 0$	Maccheroni et al. (2006)

Smooth ambiguity preferences (Klibanoff et al., 2005) can also be written in the form (1). Under these preferences, the DM evaluates σ by $\int \varphi\big(u(\sigma, q)\big)\mu(\mathrm{d}q)$, where μ is a second-order prior over $\Delta(\Omega)$ and φ is an increasing and concave function. To construct the ambiguity operator G, define $I(z) = \varphi^{-1}\big(\int \varphi(q \cdot z)\mu(\mathrm{d}q)\big)$ for $z \in \mathbb{R}^\Omega$. Then, by Cerreia-Vioglio et al. (2011), these preferences have the form (1) with

$$G(t, p, u) = \sup\big\{I(z) : z \in \mathbb{R}^\Omega,\ p \cdot z \leq t\big\}, \quad t \in \mathbb{R},\ p \in P,\ u \in \bar{\mathcal{U}}, \tag{2}$$

where P is the closed convex hull of the support of μ.

Smooth ambiguity with arbitrary φ and μ is too general to make any specific conclusions, so let us consider a special case, *exponential smooth ambiguity*, with $\varphi(t) = -e^{-t/\theta}$, where $\theta > 0$ and μ assigns a positive probability $\bar{p}(\omega) > 0$ to each degenerate

prior in $\Delta(\Omega)$. Then $P = \Delta(\Omega)$, and the ambiguity operator has the closed form with a relative entropy shift, $G(t, p, u) = t + \theta R(p\|\bar{p})$, so this is the multiplier preferences criterion.

3. Consistency and Falsifiability

We now introduce two desirable properties of the ambiguity operator G: consistency and falsifiability. We then show how these criteria apply to the ambiguity operators in Table 1.

We first define the equivalence class of utility functions that represent the same preferences as u under expected utility. Say that utility functions u and v represent the same preferences if both compare actions in the same way in every state, up to a scaling constant. This is formalized in the following definition.

Definition 1. Two utility functions $u, v \in \bar{\mathcal{U}}$ *represent the same preferences under expected utility* if there exists $c > 0$ such that

$$v(x, \omega) - v(y, \omega) = c(u(x, \omega) - u(y, \omega))$$

for all $(x, y) \in X \times X$ and all $\omega \in \Omega$.

For every $u \in \bar{\mathcal{U}}$, let $\mathcal{V}_u$ be the equivalence class of the utility functions that represent the same preferences as u under expected utility.

3.1. Consistency.

Informally speaking, consistency is a strong property implying that the optimal choices under ambiguity do not depend on which primitive utility function in the equivalence class $\mathcal{V}_u$ is used to compare actions under criterion (1).

Analogously to the equivalence under expected utility (Definition 1), we say that two utility functions represent the same preferences under ambiguity with operator G if the DM compares (mixed) actions in the same way, up to a scaling constant, for every subset of priors. To state this formally, let us denote by $W_G(\sigma; P, v)$ the numerical evaluation of a mixed action σ under criterion (1) for a given set of priors P and a given utility function v:

$$W_G(\sigma; P, v) = \min_{p \in P} G(v(\sigma, p), p, v). \tag{3}$$

Definition 2. Two utility functions $u, v \in \bar{\mathcal{U}}$ *represent the same preferences under ambiguity with operator* G if there exists $c > 0$ such that

$$W_G(\sigma; P, v) - W_G(\sigma'; P, v) = c\,(W_G(\sigma; P, u) - W_G(\sigma'; P, u)) \tag{4}$$

for all $(\sigma, \sigma') \in \Delta(X) \times \Delta(X)$ and all $P \in \mathcal{P}$.

Consistency is a requirement on behavior: for every set of priors, the DM's preferences over mixed actions under ambiguity do not depend on which utility function in the equivalence class is used to represent her fundamental preferences in the expected utility framework.

Definition 3. An ambiguity operator G is *consistent* if every two utility functions that represent the same preferences under expected utility also represent the same preferences under ambiguity with operator G.

To simplify notation, let $G_v(\sigma, p)$ be the operator G evaluated at $v \in \bar{\mathcal{U}}$, so

$$G_v(\sigma, p) = G(v(\sigma, p), p, v).$$

Proposition 1. *An ambiguity operator G is consistent if and only if for every $u \in \bar{\mathcal{U}}$ the following holds. For every $v \in \mathcal{V}_u$, there exist $b_v > 0$ and $\kappa_v \colon \Delta(\Omega) \to \mathbb{R}$ such that*

$$G_v(\sigma, p) = b_v\big(G_u(\sigma, p) + \kappa_v(p)\big) \quad \textit{for all } (\sigma, p) \in \Delta(X) \times \Delta(\Omega) \tag{5}$$

and, for each pair $(p, q) \in \Delta(\Omega) \times \Delta(\Omega)$, the expression

$$\big|G_u(\sigma, p) - G_u(\sigma, q) + \kappa_v(p) - \kappa_v(q)\big| - \big|G_u(\sigma, p) - G_u(\sigma, q)\big| \tag{6}$$

is constant in σ.

The proof is in the Appendix. Three features of Proposition 1 are worth noting. First, even though the optimization criterion W_G under an operator G relies on minimization over arbitrary sets P of priors, to determine whether G is consistent, it is sufficient to consider only pairs of priors. Second, condition (5) shows that, for a consistent operator G, G_v must be an affine transformation of G_u for any u and v in the same equivalence class, with a scale b_v common to all priors and a shift $\kappa_v(p)$ that may depend on the prior. Operators whose slope in t changes with the utility

representation differently across priors, such as the competitive ratio, fail this condition, thus failing consistency. Third, observe that expression (6) is constant in σ if and only if at least one of the following three conditions holds:

(a) $\kappa_v(p) = \kappa_v(q)$, in which case (6) is identically zero;

(b) p dominates q under both u and v, in the sense that $G_u(\sigma, p) \geq G_u(\sigma, q)$ and $G_u(\sigma, p) + \kappa_v(p) \geq G_u(\sigma, q) + \kappa_v(q)$ for all σ (or, symmetrically, q dominates p under both u and v), in which case (6) depends only on $\kappa_v(p) - \kappa_v(q)$;

(c) $G_u(\sigma, p) - G_u(\sigma, q)$ does not depend on σ, in which case (6) is, obviously, constant in σ.

These conditions are distinct and do not nest one another. For example, minmax regret satisfies (a) for every pair of priors, but both (b) and (c) fail whenever the sets of best actions under the two priors are disjoint. Moreover, none of these conditions is necessary on its own. A consistent operator may satisfy different conditions for different pairs of priors, as long as one of these conditions is satisfied for each pair.

3.2. **Falsifiability.** In contrast to consistency, falsifiability is a weak property, only implying that "not everything goes": given a hypothesis about the DM's preferences (represented by an equivalence class $\mathcal{V}_u$) and a subset of priors, there exists an undominated choice of the DM that is not optimal under the hypothesis, and thus, if observed, proves the hypothesis to be wrong.

Definition 4. An action $x \in X$ is *undominated under* $(u, P) \in \bar{\mathcal{U}} \times \mathcal{P}$ if there does not exist a mixed action $\sigma \in \Delta(X)$ such that $u(x, p) \leq u(\sigma, p)$ for all $p \in P$, with strict inequality for some $p \in P$.

Given a utility function u and a set of priors P, say that an ambiguity operator G is unfalsifiable at $(\mathcal{V}_u, P)$ if, for every undominated action under (u, P), there exists a utility function v in the equivalence class $\mathcal{V}_u$ such that this action is optimal for the DM. In other words, there is no action the observation of which would falsify the hypothesis that the DM's preferences are given by the equivalence class $\mathcal{V}_u$. Furthermore, if the above holds for all equivalence classes and all sets of priors, then we call G unfalsifiable.

Definition 5. An ambiguity operator G is *unfalsifiable* if, for every $u \in \bar{\mathcal{U}}$, every $P \in \mathcal{P}$, and every action $x \in X$ that is undominated under (u, P), there exists $v \in \mathcal{V}_u$

such that

$$x \in \arg\max_{\sigma \in \Delta(X)} W_G(\sigma; P, v).$$

Proposition 2. *If an ambiguity operator G satisfies $G(\cdot, \cdot, v) = G(\cdot, \cdot, u)$ for all $u \in \bar{\mathcal{U}}$ and all $v \in \mathcal{V}_u$, then it is unfalsifiable.*

The proof is in the Appendix.

Remark 1. Every consistent operator is falsifiable.[2] To see why, consider u for which two actions x and x' are uniquely best in different states, and let P consist of the two corresponding degenerate priors. Let x be optimal under P. If the operator is unfalsifiable, then x' can also be made optimal by a choice of a different utility function in the equivalence class $\mathcal{V}_u$. But then the comparison between these actions is not the same under two utility functions in the same equivalence class, so the operator is not consistent.

TABLE 2. Consistency and falsifiability of preferences under ambiguity.

Criterion	Consistent	Falsifiable
Maxmin utility	No	No
Minmax regret	Yes	Yes
Maxmin safety	Yes	Yes
Competitive ratio	No	Yes
Normalized maxmin utility	Yes	Yes iff $\lvert X\rvert \geq 3$
Reference-based utility	Yes	Yes
Multiplier preferences	No	No
Variational preferences	No	No

3.3. **Application to known preferences under ambiguity.** We now apply Propositions 1 and 2 to the ambiguity operators in Table 1. By Proposition 2, maxmin utility, multiplier preferences, and variational preferences are unfalsifiable (hence, inconsistent), since their operators $G(t, p, u)$ do not depend on the utility function u. Next, by Proposition 1, minmax regret, maxmin safety, reference-based utility, and normalized maxmin utility are consistent (hence, falsifiable). In particular, each of

[2]For falsifiability of normalized maxmin utility, we need at least three actions, $\lvert X\rvert \geq 3$, as if there are only two actions, the normalization assumption leaves only one undominated action, so it is trivially unfalsifiable.

them satisfies condition (a), i.e., $\kappa_v(p) = \kappa_v(q)$, for every pair of priors (p, q). Finally, the competitive ratio is falsifiable but inconsistent, which is shown directly in Appendix A.3. These results are summarized in Table 2.

Appendix A. Proofs

By Definition 1, $v \in \bar{\mathcal{U}}$ belongs to $\mathcal{V}_u$ if and only if, for some $c > 0$ and $\lambda\colon \Omega \to \mathbb{R}$,

$$v(x, \omega) = cu(x, \omega) + \lambda(\omega) \quad \text{for all } x \in X \text{ and } \omega \in \Omega. \tag{7}$$

Indeed, the definition says that $v(x, \omega) - cu(x, \omega)$ is independent of x. By linearity, $v(\sigma, p) = cu(\sigma, p) + \lambda(p)$, where $\lambda(p) = \sum_\omega p(\omega)\lambda(\omega)$.

A.1. **Proof of Proposition 1.** Fix $u \in \bar{\mathcal{U}}$ and $v \in \mathcal{V}_u$. For $b_v > 0$, write

$$D_P(\sigma) = W_G(\sigma; P, v)/b_v - W_G(\sigma; P, u).$$

By (4), u and v represent the same preferences under ambiguity if and only if some $b_v > 0$ makes D_P constant in σ for every P.

For $P = \{p\}$, constancy means precisely (5), with $\kappa_v(p) = D_{\{p\}}(\sigma)$. Given (5), twice $D_{\{p,q\}}(\sigma)$ equals $\kappa_v(p) + \kappa_v(q)$ minus expression (6), since $2\min\{a, b\} = a + b - |a - b|$. Thus the two conditions in the proposition are necessary.

For sufficiency, suppose these conditions hold. First take a finite P. For each σ, choose minimizers $p, q \in P$ under u, v, respectively. Restricting P to $\{p, q\}$ preserves both minima, so $D_P(\sigma) = D_{\{p,q\}}(\sigma)$. Each pair contributes only one value, since $D_{\{p,q\}}$ is constant. Thus D_P takes finitely many values. It is continuous on the connected set $\Delta(X)$, so it is constant.

Now take any compact P and any σ, σ'. By continuity in the prior, the minima for these two actions under u and v are attained. Let P_0 consist of one minimizer for each of the four problems. Then $D_P(\sigma) = D_{P_0}(\sigma) = D_{P_0}(\sigma') = D_P(\sigma')$, where the middle equality follows from the finite case. This proves sufficiency.

To verify the characterization in (a)–(c), write $d = G_u(\sigma, p) - G_u(\sigma, q)$ and $k = \kappa_v(p) - \kappa_v(q)$. If $k = 0$, (a) holds and (6) vanishes. Otherwise interchange p and q if needed so that $k > 0$. The function $d \mapsto |d + k| - |d|$ equals $-k$ on $(-\infty, -k]$, increases strictly on $[-k, 0]$, and equals k on $[0, \infty)$. Thus it is constant over the

realized values of d exactly when they all lie in one of the two outer intervals, as in (b), or d itself is constant, as in (c). □

A.2. **Proof of Proposition 2.** Fix u, P, and an undominated action x. Write $G(t,p)$ for the common operator on $\mathcal{V}_u$, and let K be the convex hull of P.

Step 1. We show that x is a best action under some $p_0 \in \operatorname{ri} K$. Let L be the linear span of K. For each $y \in X$, let d_y be the linear function on L given by $d_y(z) = \sum_\omega (u(y,\omega) - u(x,\omega))z(\omega)$. Suppose the cone $C = \{z \in L : d_y(z) \leq 0 \text{ for all } y\}$ misses $\operatorname{ri} K$. Then it also misses the cone $O = \{\alpha p : \alpha > 0,\ p \in \operatorname{ri} K\}$, which is open in L since every prior sums to one. Separating these cones in L gives a linear function f with $f \leq 0$ on C and $f > 0$ on O. By Farkas' lemma, $f = \sum_y \alpha_y d_y$ for some $\alpha_y \geq 0$. By continuity, $f \geq 0$ on K, and f must be positive somewhere on P, since otherwise it vanishes on its convex hull. Normalizing the coefficients α_y to sum to one therefore gives a mixed action that dominates x, a contradiction.

Step 2. We construct an affine function a that majorizes the zero of $G(\cdot, p)$ on P and equals it at priors whose convex combination is p_0. Let $r(p)$ be this unique zero, which exists by continuity, strict monotonicity, and surjectivity. For any $\varepsilon > 0$, $G(r(p) - \varepsilon, q) < 0 < G(r(p) + \varepsilon, q)$ for q near p, by continuity. Thus $|r(q) - r(p)| < \varepsilon$, so r is continuous.

The convex hull H of $\{(p, r(p)) : p \in P\}$ is compact. Its upper boundary $\hat{r}(q) = \max\{s : (q,s) \in H\}$ is concave on K by convexity of H. Since $p_0 \in \operatorname{ri} K$, a supergradient gives an affine function a on the affine hull of K with $a \geq \hat{r}$ on K and $a(p_0) = \hat{r}(p_0)$. Extend a affinely to the simplex. Write $(p_0, \hat{r}(p_0))$ as a convex combination of $(p_i, r(p_i))$ with $p_i \in P$ and positive weights μ_i. Since $a \geq r$ on P and $\sum_i \mu_i (a(p_i) - r(p_i)) = a(p_0) - \hat{r}(p_0) = 0$, we have $a(p_i) = r(p_i)$ for every i.

Step 3. Set $v(y,\omega) = u(y,\omega) - u(x,\omega) + a(\delta_\omega)$. Then $v \in \mathcal{V}_u$ by (7). Since a is affine, $v(x,p) = a(p) \geq r(p)$ on P, with equality at each p_i. Thus $W_G(x; P, v) = 0$. For any mixed action σ,

$$\sum_i \mu_i \big(v(\sigma, p_i) - v(x, p_i)\big) = u(\sigma, p_0) - u(x, p_0) \leq 0.$$

The equality follows from $\sum_i \mu_i p_i = p_0$, and the inequality holds because x is best under p_0. At some p_i, then, $v(\sigma, p_i) \leq v(x, p_i) = r(p_i)$. By monotonicity, $W_G(\sigma; P, v) \leq G(r(p_i), p_i) = 0$. Thus x is optimal under v. □

A.3. **Entries of Table 2.** *Maxmin utility, multiplier preferences, and variational preferences.* In each case $G(t,p,u) = t + \phi(p)$, where ϕ is zero for maxmin utility and is the cost of the prior for the other two criteria. Since G is independent of u, Proposition 2 gives unfalsifiability. Remark 1 then gives inconsistency.

Minmax regret, maxmin safety, and reference-based utility. Under $v = cu + \lambda$, each benchmark changes by the same scale and shift as expected utility. Thus $G_v = cG_u$ in all three cases. Proposition 1 gives consistency and Remark 1 gives falsifiability.

Normalized maxmin utility. Both numerator and denominator scale by c, so $G_v = G_u$ and the criterion is consistent. With two actions, their expected-utility difference is affine in p and never zero, since $\bar{u}(p) > \underline{u}(p)$ throughout the simplex. Thus the same action is strictly best at every prior. It is the only undominated action and is always optimal, so the criterion is unfalsifiable.

With at least three actions, give x and y payoffs $(2,1)$ and $(1,2)$ in two states, give all other actions zero, and duplicate the first state if there are others. The denominator is positive at every prior. Let P consist of the degenerate priors at these two states. Action x is undominated because it is uniquely best in the first state. Its worst-case normalized utility is $1/2$, whereas the equal mixture of x and y yields $3/4$ in both states. Since $G_v = G_u$, observing x falsifies the hypothesis $\mathcal{V}_u$.

Competitive ratio. Fix two actions x, y and two states ω_1, ω_2. Give x payoff 2 in ω_1 and 1 elsewhere, and y payoff 2 in ω_2 and 1 elsewhere. Give all other actions payoff zero. Then $\bar{u}(p) \geq 1$ for every prior. Let $p_i = \delta_{\omega_i}$ and $P = \{p_1, p_2\}$.

For inconsistency, add 1 to every payoff in ω_2 to obtain $v \in \mathcal{V}_u$. The absolute gap between the ratios of x and y stays at $1/2$ in ω_1, so (5) requires $b_v = 1$. In ω_2 the gap falls from $1/2$ to $1/3$, which requires $b_v = 2/3$. Thus the competitive ratio is inconsistent by Proposition 1.

For falsifiability, x is undominated because it is uniquely best under p_1. For every admissible $v \in \mathcal{V}_u$, $G_v(x, p_1) = 1$ and $G_v(x, p_2) < 1$, since y is strictly better under p_2. Mixing a small positive probability of y with x strictly raises the ratio under p_2. By continuity, the ratio under p_1 also stays above $G_v(x, p_2)$. Thus the mixture strictly improves on x under P, so no admissible v makes x optimal. □

References

Anscombe, F. J. and R. J. Aumann (1963): "A definition of subjective probability," *Annals of Mathematical Statistics*, 34, 199–205.

Bewley, T. F. (1986): "Knightian Decision Theory: Part 1," Cowles Foundation Discussion Paper 807, Cowles Foundation, Yale University.

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

Chambers, C. P., F. Echenique, and K. Saito (2016): "Testing Theories of Financial Decision Making," *Proceedings of the National Academy of Sciences*, 113, 4003–4008.

Dhillon, A. and J.-F. Mertens (1999): "Relative Utilitarianism," *Econometrica*, 67, 471–498.

Echenique, F. (2020): "New Developments in Revealed Preference Theory: Decisions Under Risk, Uncertainty, and Intertemporal Choice," *Annual Review of Economics*, 12, 299–316.

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

Gulko, B. and S. Leung (2013): "Maximin Safety: When Failing to Lose Is Preferable to Trying to Win," in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2013)*, Springer, Lecture Notes in Computer Science, 254–265.

Hansen, L. P. and T. J. Sargent (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91, 60–66.

Hayashi, T. (2008): "Regret Aversion and Opportunity Dependence," *Journal of Economic Theory*, 139, 242–268.

Karni, E. (1985): *Decision Making under Uncertainty: The Case of State-Dependent Preferences*, Harvard University Press, Cambridge, MA.

Karni, E., D. Schmeidler, and K. Vind (1983): "On state dependent preferences and subjective probabilities," *Econometrica*, 51, 1021–1031.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73, 1849–1892.

Luce, R. D. and H. Raiffa (1957): *Games and Decisions: Introduction and Critical Survey*, Wiley, New York.

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

Milnor, J. (1954): "Games against nature," in *Decision Processes*, ed. by R. M. Thrall, C. H. Coombs, and R. L. Davis, Wiley, New York, 49–59.

Moulin, H. and J. P. Vial (1978): "Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon," *International Journal of Game Theory*, 7, 201–221.

Ortoleva, P. (2010): "Status Quo Bias, Multiple Priors and Uncertainty Aversion," *Games and Economic Behavior*, 69, 411–424.

Savage, L. J. (1951): "The theory of statistical decision," *Journal of the American Statistical Association*, 46, 55–67.

Sleator, D. and R. Tarjan (1985): "Amortized efficiency of list update and paging rules," *Communications of the ACM*, 28, 202–208.

Stoye, J. (2011): "Axioms for Minimax Regret Choice Correspondences," *Journal of Economic Theory*, 146, 2226–2251.