Worst- and Best-Case Expected Utility and Ordinal Meta-Utility*

Simon Grant   Patricia Rich   Jack Stecher

April 6, 2021

Abstract

We provide a theory of decision under ambiguity in which the worst-case and best-case expected utility representations serve as sufficient statistics for a decision maker's preferences. A second-order, ordinal "meta-utility" function aggregates these lower and upper expected utilities. This meta-utility function is just an instance of the familiar ordinal utility functions from intermediate microeconomics. We show that several important models of decision under ambiguity are special cases of constant elasticity of substitution meta-utilities, and provide further generalizations. The key insight is that ambiguity attitude is simply the marginal rate of substitution in the meta-utility between worst-case and best-case expected utilities.

JEL Classification: D80, D81

Key words: ambiguity, expected utility, optimism, pessimism, second-order utility.

Wanting you to order his new creation, a chef might describe it in terms of more familiar foods, thereby making it seem less exotic than you first feared. In this article, we hope to do something similar for the theory of decision under ambiguity, presenting a theory which subsumes several existing models and sheds light on others. We show that our general theory follows a recipe that uses two familiar ingredients from intermediate microeconomics. If you hesitate to order pheasant, we want to reassure you that it tastes like chicken.

Our two familiar ingredients are ordinal utility theory and expected utility. With no uncertainty, the standard ordinal utility theory we learn in intermediate microeconomics is the benchmark. We might use different variations for different purposes (as a chef uses different herbs), but ordinal utility is the base, with well-understood properties. With unique probabilities, expected utility is the default normative theory. Again, we may find it helpful to toss in chervil or tarragon, but expected utility gives us a general framework with a well-understood underlying structure.

The theory we develop here uses these two ingredients, expected utility and ordinal utility theory, and shows that combining them gives us a general way to study decision under ambiguity. We use expected utility first. In situations of ambiguity, as in Ellsberg's (1961) classic example, probabilities are not measurable with respect to an agent's information partition. Nevertheless, an outer and inner measure of any event is always available, so at least two expected utilities characterize each ambiguous act: there is a worst-case and a best-case expected utility. Label them $\mathbf{w}$ and $\mathbf{b}$, respectively.[1]

These expected utilities serve as "goods" for our second ingredient, a two-good ordinal utility function. In intermediate microeconomics, the ordinal utility takes wine and beer as inputs, and expresses how a decision maker (hereafter DM) values the trade-offs between them. In what follows, the ordinal utility operates at a meta-level, taking $\mathbf{w}$ and $\mathbf{b}$ as the inputs. Mixing these ingredients in this order is the recipe for what we call *ordinal Hurwicz expected utility* (OHEU), a theory that generalizes Hurwicz (1951) by finding a general, not necessarily linear, way to combine $\mathbf{w}$ and $\mathbf{b}$.[2]

We provide an axiomatization and representation theorem for an OHEU, and a fundamental insight that follows from our representation: in our framework and all those it subsumes, ambiguity attitude is just the marginal rate of substitution (MRS) between the best- and worst-case expected utilities in the meta-utility function.

---

[1]This measure-theoretic approach is pursued in numerous theories of decision under ambiguity; examples include Good (1966, 1983), Suppes (1974), and Binmore (2009). In the special case where the inner probability of an event is zero and the outer probability is one, the decision maker is completely ignorant in the sense of Arrow and Hurwicz (1972); see also the related random-set approach discussed in Pintér (2019) and citations therein. A benefit of this approach is that it captures non-probabilistic information through belief functions (see Jaffray and Wakker, 1993). As we detail below, our approach provides insights into several other models that can depend on more than $\mathbf{w}$ and $\mathbf{b}$. Moreover, in these other models, there are often important special cases in which $\mathbf{w}$ and $\mathbf{b}$ are sufficient statistics for the decision maker's preferences (see also Chateauneuf, Eichberger and Grant, 2007; Grant and Polak, 2013).

[2]Gul and Pesendorfer (2014) go in the opposite direction, applying what we interpret as a cardinal utility state-by-state to a best and worst outcome, then aggregating. Their approach produces a linear averaging; see Gul and Pesendorfer (2015).

The meta-utility approach's ability to provide a deeper, systematic understanding of ambiguity is a recurring theme in this paper. It immediately makes clear how to interpret a number of existing models, extend them, and unify their seemingly disparate notions of ambiguity attitude. This is because meta-utility functions are implicit throughout the literature. Textbook ordinal utility examples turn out to capture several prominent models of decision under ambiguity (and special cases of some others).

To illustrate, consider three well-known models that depend only on $\mathbf{w}$ and $\mathbf{b}$. The maxmin expected utility (MEU) model of Wald (1950) and Gilboa and Schmeidler (1989) views a DM as pessimistic, maximizing the worst-case expected utility representation, so that $U(\mathbf{w}, \mathbf{b}) = \min(\mathbf{w}, \mathbf{b}) = \mathbf{w}$. The α-maxmin expected utility (α-MEU) model (Hurwicz, 1951; Gul and Pesendorfer, 2015), with $U(\mathbf{w}, \mathbf{b}) = \alpha\mathbf{w} + (1-\alpha)\mathbf{b}$, views a DM as calculating a weighted average of $\mathbf{w}$ and $\mathbf{b}$, with weight $\alpha \in [0, 1]$ on $\mathbf{w}$ reflecting the DM's degree of pessimism. The geometric α-MEU model (Binmore, 2009) instead considers a geometric weighted average, with $U(\mathbf{w}, \mathbf{b}) = \mathbf{w}^\alpha \mathbf{b}^{1-\alpha}$.

All of these models are instances of constant elasticity of substitution meta-utility functions, corresponding to perfect complements, perfect substitutes, and Cobb-Douglas. We can view them as variations of one theory, rather than as competitors, expressing them as follows:

$$U(\mathbf{w}, \mathbf{b}) = [\alpha\mathbf{w}^\rho + (1 - \alpha)\mathbf{b}^\rho]^{1/\rho}.$$

The meta-utility perspective makes salient—and easy to visualize—the different hypotheses underlying these models. When there exist a best and worst outcome, we can normalize their utilities to 1 and 0 respectively, and display the commodity space as the upper triangle of the unit square (see Figures 1, 2, and 3). As $\mathbf{w}$ is never greater than $\mathbf{b}$, the figures show the portion of the indifference curves below the 45° line as dotted lines or curves.

Our chef from above hopes to give you a way to appreciate his new creation and maybe even serve it to your own guests, not just taste it once. Similarly, we want to introduce you to the meta-utility approach for more than its own sake; it has much to teach us. A primary benefit of the approach is that it clarifies the concept of ambiguity attitude, which the literature views as an open problem. For example, Ahn et al.'s (2014) experiments treat ambiguity attitude in a model-dependent way, defined in terms of the pessimism parameter α in the α-MEU model and in a somewhat more complex way in the smooth model (Klibanoff, Marinacci and Mukerji, 2005). There is no consensus on a general definition. We show that the widely used definitions have more in common than is initially apparent, however: all are closely connected to the MRS between $\mathbf{w}$ and $\mathbf{b}$, a relationship the meta-utility systematization immediately reveals.

As a general definition of ambiguity attitude, the MRS captures a DM's attitude toward increases in utility dispersion—that is, to midpoint-preserving spreads in the range of expected utilities. We
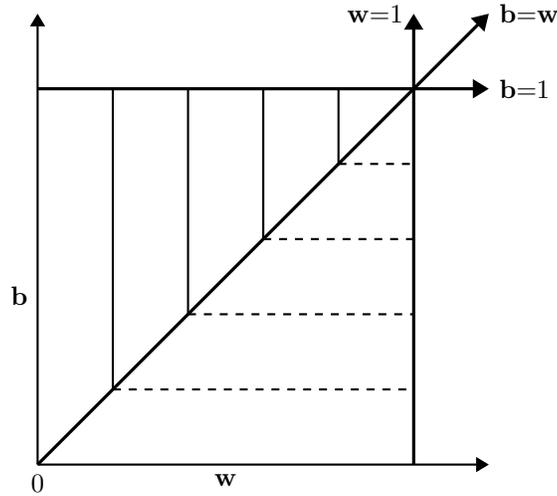
Figure 1: Leontief Meta-utilities (Gilboa-Schmeidler)

show that an agent is averse to midpoint-preserving spreads if and only if the MRS in the meta-utility function is above one. For a DM with α-MEU preference, the MRS is above one if and only if α is above one-half, the standard notion of ambiguity aversion for the α-MEU model.

Our main insights about ambiguity attitude carry over to models which approach ambiguity from a different direction, such as a failure to reduce compound lotteries, rather than from non-measurability. We show that our MRS condition has a close relationship with the measures of ambiguity attitude in the smooth model of Klibanoff, Marinacci and Mukerji (2005) and in the variational model of Maccheroni, Marinacci and Rustichini (2006). Similar reasoning applies to related models, such as Siniscalchi (2009); Seo (2009). The interpretations of the sources of these preferences differ across the models; by highlighting these differences, we provide new insights on the underlying preference structures these models aim to capture.

Our ability to subsume many existing models has important advantages for applications: it enables one to simultaneously consider agents who treat ambiguity differently, while cleanly separating the effects of the shape of the meta-utility function, of ambiguity attitude, and of ambiguity itself. We provide a new insight about these distinct effects below, in a setting of risky debt with possible strategic default and costly verification. As it turns out, ambiguity can reduce the probability of strategic default even if the parties to a debt contract are ambiguity neutral. That is, in strategic interaction, ambiguity neutrality does not imply insensitivity to ambiguity.

The rest of this article is organized as follows: Section 1 gives preliminary definitions. Section 2 provides our axioms and main representation result. Section 3 illustrates how the MRS can be viewed as a general definition of ambiguity attitude. Section 4 applies our approach to a strategic setting.
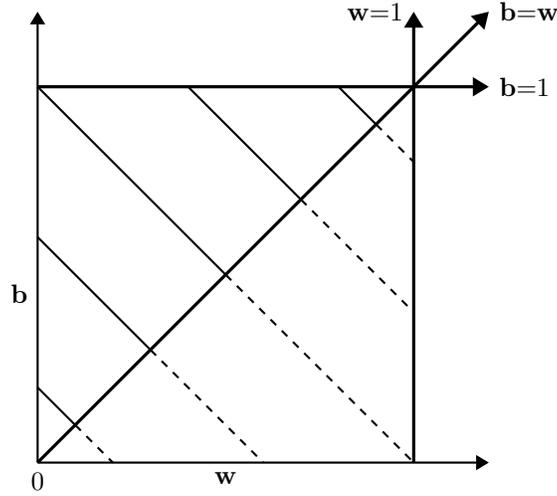
3

Figure 2: Perfect Substitutes Meta-utilities (Hurwicz)

1 Preliminaries

We use a version of Savage's (1954) setting of purely subjective uncertainty, similar to Grant, Rich and Stecher (2020). A companion paper (Grant, Rich and Stecher, 2021) extends the analysis presented here to the Anscombe and Aumann (1963) framework. The set of final outcomes is a non-degenerate interval $X = [\ell, m] \subset \mathbb{R}$. We assume uncertainty is described by a set of states, Ω, and a σ-algebra, Σ, of subsets of Ω (called *events*). For events A, B, $B \backslash A$ denotes the relative complement of A with respect to B. The objects of choice, denoted by $\mathcal{F}$, are simple acts, that is, mappings $f \colon \Omega \to X$ that are Σ-measurable and have finite range. We identify any $x \in X$ with the constant act $f = x$. For any pair of acts $f, g \in \mathcal{F}$ and any event $A \in \Sigma$, we write $f_A g$ for the act that agrees with f on A and with g on $\Omega \backslash A$. The binary relation $\succsim$ on $\mathcal{F}$ denotes the DM's preferences over acts.

In addition to the primitives Ω, X, Σ, and $\succsim$, we shall fix a sub-σ-algebra $\mathcal{R} \subset \Sigma$ to denote the set of *risky* events. These are the events to which the DM assigns a unique probability.[3] A subset of the risky events are the null events to which the DM can assign zero weight. Formally, an event N is *null* if $f \sim g_N f$ for all $f, g \in \mathcal{F}$. Let $\mathcal{N}$ denote the set of null events. We call an act f *risky* if it is $\mathcal{R}$-measurable. Let $\mathcal{G} \subset \mathcal{F}$ denote the set of risky acts.

---

[3] For clarity of exposition, we assume that $\mathcal{R}$ is a σ-algebra. For extensions to more general settings, see Grant, Rich and Stecher (2020).
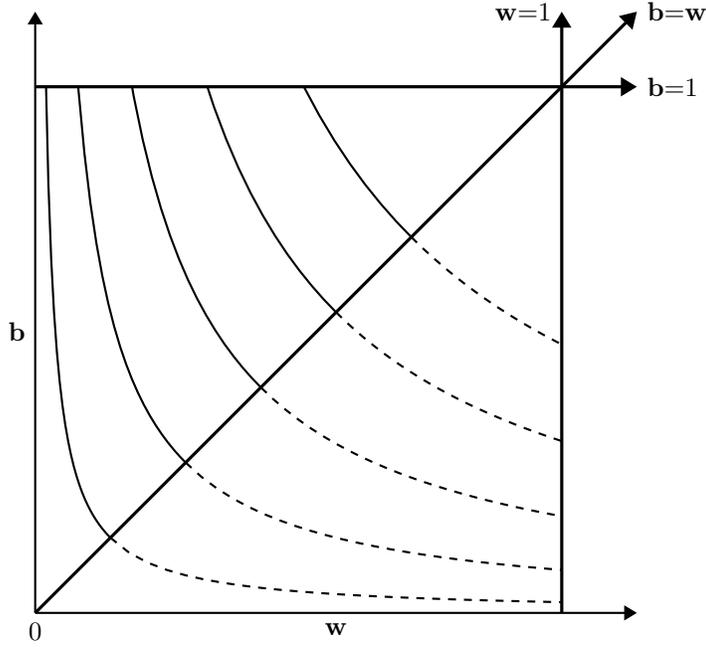
Figure 3: Cobb-Douglas Meta-utilities (Binmore)

2   Characterization

We begin our characterization of the family of preferences that admit Ordinal Hurwicz Expected Utility (OHEU) representations with the definition of a prior.

Definition 1 (Prior) *A prior π is a countably-additive and convex-ranged probability measure defined on the set of risky events $\mathcal{R}$.*

Countable additivity is defined in the usual way. Convex-ranged means that for any $A \in \mathcal{R}$ and any $r \in (0,1)$ there exists $B \subset A$, $B \in \mathcal{R}$ for which $\pi(B) = r\pi(A)$.

If we associate a prior π with a DM's preference relation $\succsim$, then we may view the restriction of $\succsim$ to $\mathcal{G}$ as reflecting the DM's risk preferences since the prior π can be used to map each (risky) act in $\mathcal{G}$ to a lottery as follows: Let $\Delta_0(X)$ denote the set of simple lotteries defined on X. That is, each $L \in \Delta_0(X)$ has finite support which means we can identify it as a function $L : X \to [0,1]$, satisfying $\sum_{x \in X} L(x) = 1$. For each $f \in \mathcal{G}$, we identify $\pi \circ f^{-1}$ with the lottery $L \in \Delta_0(X)$ in which $L(x) = \pi \circ f^{-1}(x)$.

In order for the restriction of $\succsim$ to $\mathcal{G}$ to fully characterize risk preferences, the range of the mapping $f \to \pi \circ f^{-1}$ should be $\Delta_0(X)$. That this is the case is shown in Grant, Rich and Stecher (2020, Lemma 2).

The standard model for risk preferences—(subjective) expected utility—uses a Bernoulli utility index $v: X \to \mathbb{R}$ in addition to a prior. We therefore define:

Definition 2 (Subjective Expected Utility Maximization) *The restriction of the preference relation $\succsim$ to risky acts $\mathcal{G}$ conforms to* subjective expected utility maximization *if there exists a prior π and a Bernoulli utility index v such that for every pair of acts f and g in $\mathcal{G}$,*

$$f \succsim g \iff \sum_{x \in X} \pi\left(f^{-1}(x)\right) v(x) \geqslant \sum_{x \in X} \pi\left(g^{-1}(x)\right) v(x).$$

We take subjective expected utility maximization as our starting point and consider the implications of additional properties of the preference relation $\succsim$ for its representation on the entire domain $\mathcal{F}$. In order to evaluate a non-risky act $f \in \mathcal{F} \backslash \mathcal{G}$, we assume that the DM constructs a greatest lower-bound expected utility from the set of risky acts that f dominates and a least upper-bound expected utility from the set of risky acts that dominate f. The DM then employs a meta-utility aggregator (defined below) to combine these two extreme expected utilities into a single number by which f is then ranked against other acts.

Definition 3 (Meta-utility Aggregator) *A* meta-utility aggregator *is a continuous increasing bivariate onto function*

$$U: \{(u_1, u_2) \in v(X) \times v(X): u_1 \leqslant u_2\} \to v(X),$$

in which $U(u, u) = u$, for all $u \in v(X)$.

Without this last requirement, U would be a purely ordinal function; the requirement that $U(u, u) = u$ selects a unique function that makes the meta-utility agree with expected utility when the latter is uniquely defined.

We can now state the definition of an ordinal Hurwicz expected utility.

Definition 4 (OHEU Representation) *The function $W: \mathcal{F} \to \mathbb{R}$ is an* Ordinal Hurwicz Expected Utility *(OHEU) if there is a triple $\langle \pi, v, U \rangle$, where π is a prior, v is a utility index, and U is a meta-utility aggregator such that*

$$W(f) = U\left(\underline{v}_\pi^f, \bar{v}_\pi^f\right)$$

where

$$\underline{v}_\pi^f = \sup_{g \in \mathcal{G} \,:\, g \leqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x)\,,$$

$$\bar{v}_\pi^f = \inf_{g \in \mathcal{G} \,:\, g \geqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x)\,.$$

We note that Definition 4 requires an infinite state space, as we require our prior to be convex-ranged. It is, however, straightforward to extend to a finite state space, following the approach in Gul and Pesendorfer (2014) and Grant, Rich and Stecher (2020).

Turning now to the characterization, we begin with three standard axioms.

Axiom 1 (Ordering) *The binary relation $\succsim$ is complete and transitive.*

Axiom 2 (Monotonicity) *For any pair of acts $f, g \in \mathcal{F}$, if $f > g$, then $f \succ g$.*

Axiom 3 (Continuity) *For any sequence of pairs of acts $\{f_n, h_n\}_{n=1}^\infty$ with $f_n \succsim h_n$ for all n, if $f = \lim_{n \to \infty} f_n$ and $h = \lim_{n \to \infty} h_n$ then $f \succsim h$.*

The fourth (and new) axiom formalizes the following intuition. In order to evaluate an arbitrary act f, the DM first considers the set of risky acts that f statewise dominates, evaluates their certainty equivalents, and takes the supremum as a lower bound for the certainty equivalent of f. Next, she considers the set of risky acts that statewise dominate f and takes the infimum of their associated certainty equivalents as an upper bound for the certainty equivalent of f. The resulting interval of outcomes embodies the ambiguity associated with evaluating the act f. The next axiom implies that these intervals fully characterize the DM's (ambiguity) preferences. Specifically, if both the lower and upper bounds for the certainty equivalent of an act f are at least as large as those for another f', then f must be weakly preferred to f'.[4]

Axiom 4 (Envelope Monotonicity) *Fix a pair of acts $f, f' \in \mathcal{F}$. If for each outcome $x \in X$:*

  I *($x \succsim g$ for all $g \in \mathcal{G}$ such that $g \leqslant f$) implies ($x \succsim g'$ for all $g' \in \mathcal{G}$ such that $g' \leqslant f'$); and,*

  II *($g' \succsim x$ for all $g' \in \mathcal{G}$ such that $g' \geqslant f'$) implies ($g \succsim x$ for all $g \in \mathcal{G}$ such that $g \geqslant f$);*

Then $f \succsim f'$.

---

[4]An alternative formulation of Axiom 4 can be stated purely in terms of acts, rather than in terms of certainty equivalents. Specifically, let $\underline{f}$ be the maximal risky act that f weakly dominates, and let $\overline{f}$ be the minimal risky act that weakly dominates f. The existence and (essential) uniqueness of these risky acts is discussed in Gul and Pesendorfer (2014, 2015). The alternative formulation of Envelope Monotonicity then states that, for $f, f' \in \mathcal{F}$, if $\underline{f} \succsim \underline{f'}$ and $\overline{f} \succsim \overline{f'}$, then $f \succsim f'$. This alternative formulation would require a slight restatement of our other axioms (details available from the authors on request). This is equivalent to the security-potential dominance axiom of Frick, Iijima and Le Yaouanq (2020); see also the discussion in Kopylov (2009).

Axioms 1–4 yield the desired representation. Note that we do not provide axioms to re-derive the expected utility part of the representation. Conditions for such a representation have been known since Savage (1954), and repeating them in this context would merely be a detour.

Theorem 1 *Suppose the restriction of the preference relation $\succsim$ to risky acts conforms to subjective expected utility maximization. Then $\succsim$ satisfies Axioms 1–4 if and only if it admits an OHEU representation.*

Furthermore, any two OHEU functions characterized by the pair of triples $\langle \pi, v, U \rangle$ and $\langle \hat{\pi}, \hat{v}, \hat{U} \rangle$ represent the same preference relation if and only if there exists $a > 0$ and $b \in \mathbb{R}$ such that $\hat{\pi} = \pi$, $\hat{v} = av + b$ and $\hat{U}(\hat{u}_1, \hat{u}_2) = aU\left(\frac{\hat{u}_1 - b}{a}, \frac{\hat{u}_2 - b}{a}\right) + b$ for all

$$(\hat{u}_1, \hat{u}_2) \in \left\{ (w, w') \in v(X) \times v(X) : \exists f \in \mathcal{F} \text{ s.t. } w = \underline{v}_\pi^f, w' = \bar{v}_\pi^f \right\}.$$

Proof. Necessity. Fix an OHEU functional $W : \mathcal{F} \to \mathbb{R}$, characterized by the triple $\langle \pi, v, U \rangle$. By construction, W generates a preference relation over acts that is complete, transitive, monotonic, and continuous, hence Axioms 1–3 are necessary. To see that Axiom 4 also holds for this preference relation, fix $f, f' \in \mathcal{F}$. If for each $x \in X$:

I $(x \succsim g$ for all $g \in \mathcal{G}$ such that $g \leqslant f)$ implies $(x \succsim g'$ for all $g' \in \mathcal{G}$ such that $g' \leqslant f')$; and,

II $(g' \succsim x$ for all $g' \in \mathcal{G}$ such that $g' \geqslant f')$ implies $(g \succsim x$ for all $g \in \mathcal{G}$ such that $g \geqslant f)$;

then we have

$$v(x) \geqslant \sup_{g \in \mathcal{G} : g \leqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x) \implies v(x) \geqslant \sup_{g \in \mathcal{G} : g \leqslant f'} \sum_{x \in X} \pi(g^{-1}(x))v(x) \text{ and}$$

$$\inf_{g \in \mathcal{G} : g \geqslant f'} \sum_{x \in X} \pi(g^{-1}(x))v(x) \geqslant v(x) \implies \inf_{g \in \mathcal{G} : g \geqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x) \geqslant v(x).$$

Hence,

$$\underline{v}_\pi^f = \sup_{g \in \mathcal{G} : g \leqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x) \geqslant \sup_{g \in \mathcal{G} : g \leqslant f'} \sum_{x \in X} \pi(g^{-1}(x))v(x) = \underline{v}_\pi^{f'} \text{ and}$$

$$\bar{v}_\pi^{f'} = \inf_{g \in \mathcal{G} : g \geqslant f'} \sum_{x \in X} \pi(g^{-1}(x))v(x) \leqslant \inf_{g \in \mathcal{G} : g \geqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x) = \bar{v}_\pi^f.$$

Since U is monotonic in both arguments, this in turn implies

$$W(f) = U\left(\underline{v}_\mu^f, \bar{v}_\mu^f\right) \geqslant U\left(\underline{v}_\mu^{f'}, \bar{v}_\mu^{f'}\right) = W(f').$$

Hence $f \succsim f'$, which is what Axiom 4 delivers. $\qquad \square$

Sufficiency. Let $\langle \pi, v \rangle$ characterize the expected representation of $\succsim$ restricted to $\mathcal{G}$, the set of risky acts. By standard arguments, Axioms 1–3 imply that $\succsim$ admits a certainty equivalent representation $c : \mathcal{F} \to X$.[5]

Next, for each act $f \in \mathcal{F}$, set

$$\underline{v}_\pi^f := \sup_{g \in \mathcal{G}\,:\, g \leqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x) \ \text{ and } \ \bar{v}_\pi^f := \inf_{g \in \mathcal{G}\,:\, g \geqslant f} \sum_{x \in X} \pi(g^{-1}(x))v(x).$$

Consider the function $\hat{U}\left(\cdot, \cdot\right) : \mathcal{I} \longrightarrow \mathbb{R}$, where

$$\mathcal{I} = \left\{ (w, w') \in v\left(X\right) \times v\left(X\right) : \exists\, f \in \mathcal{F} \text{ s.t. } w = \underline{v}_\pi^f,\, w' = \bar{v}_\pi^f \right\},$$

defined by setting $\hat{U}\left(w, w'\right) := v\left(c\left(f\right)\right)$, where $f \in \mathcal{F}$ is an act for which $w = \underline{v}_\pi^f$, $w' = \bar{v}_\pi^f$. To show $\hat{U}\left(\cdot, \cdot\right)$ is well-defined (respectively, monotonic), consider any $(w, w'), (\hat{w}, \hat{w}') \in \mathcal{I}$, with $(w, w') =$ (respectively, $\geqslant$) $(\hat{w}, \hat{w}')$. Let f and $\hat{f}$ be the two acts in $\mathcal{F}$ for which $\left(\underline{v}_\pi^f, \bar{v}_\pi^f\right) = (w, w')$ and $\left(\underline{v}_\pi^{\hat{f}}, \bar{v}_\pi^{\hat{f}}\right) = (\hat{w}, \hat{w}')$. The two (in)equalities $\underline{v}_\pi^f = (\geqslant)\, \underline{v}_\pi^{\hat{f}}$ and $\bar{v}_\pi^f = (\geqslant)\, \bar{v}_\pi^{\hat{f}}$ imply for any outcome $x \in X$:

$$v\left(x\right) \geqslant \sup_{g \in \mathcal{G}\,:\, g \leqslant f} \sum_{z \in X} \pi(g^{-1}(z))v(z) \Longleftrightarrow (\Longrightarrow)\ v\left(x\right) \geqslant \sup_{g \in \mathcal{G}\,:\, g \leqslant \hat{f}} \sum_{z \in X} \pi(g^{-1}(z))v(z) \text{ and}$$

$$\inf_{g \in \mathcal{G}\,:\, g \geqslant \hat{f}} \sum_{z \in X} \pi(g^{-1}(z))v(z) \geqslant v\left(x\right) \Longleftrightarrow (\Longrightarrow)\ \inf_{g \in \mathcal{G}\,:\, g \geqslant f} \sum_{z \in X} \pi(g^{-1}(z))v(z) \geqslant v\left(x\right).$$

Hence, we have each outcome $x \in X$:

I $\ (x \succsim g \text{ for all } g \in \mathcal{G} \text{ such that } g \leqslant f) \Longleftrightarrow (\Longrightarrow) (x \succsim g' \text{ for all } g' \in \mathcal{G} \text{ such that } g' \leqslant \hat{f})$; and,

II $\ (g' \succsim x \text{ for all } g' \in \mathcal{G} \text{ such that } g' \geqslant \hat{f}) \Longleftrightarrow (\Longrightarrow) (g \succsim x \text{ for all } g \in \mathcal{G} \text{ such that } g \geqslant f)$.

So by Axiom 4 it follows $f \sim (\succsim)\, \hat{f}$, as required.

Finally, take $U : \{(w, w') \in v\left(X\right) \times v\left(X\right) : w \leq w'\} \longrightarrow \mathbb{R}$, to be any monotonic function which agrees with $\hat{U}$ on $\mathcal{I}$. By construction the OHEU functional characterized by the triple $\langle \pi, v, U \rangle$ represents $\succsim$. ∎

3 A General Definition of Ambiguity Attitude: The MRS

3.1 Ambiguity attitude and the MRS

We now show that our approach provides a general characterization of ambiguity attitude. As a starting point, we introduce the notion of a midpoint-preserving spread:

---

[5]See, for example, the proof of Proposition 3.C.1 in Mas-Colell, Whinston and Green (1995, 47–48).

Definition 5 *Let $f, g \in \mathcal{F}$ be two acts with corresponding expected utility bounds $(\underline{v}_\pi^f, \bar{v}_\pi^f)$ and $(\underline{v}_\pi^g, \bar{v}_\pi^g)$. Say that f is a* midpoint-preserving spread *of g if, for some $\delta > 0$,*

$$\underline{v}_\pi^g - \underline{v}_\pi^f = \delta = \bar{v}_\pi^f - \bar{v}_\pi^g$$

If f is a midpoint-preserving spread of g, then we view f as more ambiguous than g.

Remark 1 *Our definition of a midpoint-preserving spread is in terms of the DM's upper and lower expected utility. In this way, we are following the tradition of Ghirardato, Maccheroni and Marinacci (2004), and focusing on ambiguity only if the DM views it as affecting his or her payoff.*

To clarify Remark 1, an act $f \in \mathcal{F} \backslash \mathcal{R}$ is ambiguous in the obvious sense of being non-risky, but the ambiguity may be inconsequential to the DM. For instance, suppose act $f \in \mathcal{F}$ depends on a two-color Ellsberg urn. If a red ball is drawn, then the DM is committed to breaking a suspect egg into a saucer, and if a black ball is drawn, the DM is committed to breaking the egg into a bowl containing other fresh eggs. Assume that the DM has a unique subjective prior over whether the egg is fresh or rotten. If each of these risky sub-acts has the same expected utility, then the ambiguity of the Ellsberg urn does not affect the DM's payoff.

For a DM with an OHEU representation, the lower and upper expected utilities are sufficient statistics for the DM's preferences. Accordingly, we say a DM exhibits *ambiguity aversion* with respect to act g if, for every midpoint-preserving spread f of g, for sufficiently small $\delta > 0$, the DM prefers g to f. Ambiguity neutrality and ambiguity affinity are defined analogously. This notion of ambiguity is consistent with Machina (2014), who argues that adding a constant to all possible expected utilities of an ambiguous act leaves the absolute ambiguity unchanged, and Izhakian (2017), who explicitly discusses spreads in the set of probabilities, and is in the spirit of Rothschild and Stiglitz's (1970) definition for risk attitude. Related definitions of increasing ambiguity are in Ghirardato and Marinacci (2002), Ghirardato, Maccheroni and Marinacci (2004), Rigotti, Shannon and Strzalecki (2008), and Ghirardato and Siniscalchi (2012).

In the remainder of this section, we explain how the MRS functions as a general definition of ambiguity attitude. The MRS in the meta-utility function $U(\mathbf{w}, \mathbf{b})$ is the standard notion of MRS, using the lower and upper expected utilities as the commodities:

$$MRS = \frac{\partial U(\mathbf{w}, \mathbf{b})/\partial \mathbf{w}}{\partial U(\mathbf{w}, \mathbf{b})/\partial \mathbf{b}}$$

Proposition 2 shows that the MRS of the meta-utility function fully characterizes attitude toward midpoint-preserving spreads:

Proposition 2 *Let $f \in \mathcal{F}$ be an act with worst-case expected utility $\boldsymbol{w}$ and best-case expected utility*

b. *A DM is averse to (respectively, is indifferent to, seeks) midpoint-preserving spreads of f if and only if the MRS of the DM's meta-utility is greater than (respectively, equal to, less than) one.*

Proof. Let $\mathbf{w}$, $\mathbf{b}$, and the meta-utility U be given. Define $c := (\mathbf{w} + \mathbf{b})/2$ as the center, i.e., midpoint of the range of expected utilities, and define $r := (\mathbf{b} - \mathbf{w})/2$ as the radius. Then

$$U(\mathbf{w}, \mathbf{b}) = U(c - r, c + r).$$

Holding c fixed, an increase in r corresponds to a midpoint-preserving spread. Using the chain rule and differentiating with respect to r,

$$\begin{aligned}
\frac{\partial U}{\partial r} &= \frac{\partial U}{\partial \mathbf{w}} \cdot \frac{\partial \mathbf{w}}{\partial r} + \frac{\partial U}{\partial \mathbf{b}} \cdot \frac{\partial \mathbf{b}}{\partial r} \\
&= \frac{\partial U}{\partial \mathbf{w}}(-1) + \frac{\partial U}{\partial \mathbf{b}}(1) \\
&= \frac{\partial U}{\partial \mathbf{b}} - \frac{\partial U}{\partial \mathbf{w}}.
\end{aligned}$$

Therefore,

$$\frac{\partial U}{\partial r} \lesseqgtr 0 \iff \frac{\partial U}{\partial \mathbf{b}} \lesseqgtr \frac{\partial U}{\partial \mathbf{w}}$$
$$\iff \frac{\partial U/\partial \mathbf{w}}{\partial U/\partial \mathbf{b}} \gtreqless 1$$

∎

Proposition 2 gives us an easy way to consider comparative ambiguity attitude. The usual definition, from Ghirardato and Marinacci (2002, Definition 4), is as follows: $\succsim_1$ is a more uncertainty averse preference than $\succsim_2$ if and only if, for all $f \in \mathcal{F}$ and for all $x \in X$, we have

$$f \succsim_1 x \Rightarrow f \succsim_2 x$$

In words, if a DM with preference $\succsim_1$ would rather absorb the ambiguity in f than have sure outcome x, then so would a DM with preference $\succsim_2$.[6]

Thanks to Proposition 2, we can extend notions of comparative ambiguity attitude to cases in which, e.g., DM1 is more ambiguity averse than DM2 relative to some acts but not to others. DM1 may have steeper indifference curves for some regions, but not others. DM1 is more ambiguity averse for act $f \in \mathcal{F}$ than DM2 is if and only if DM1 is more sensitive to changes in her worst-case expected utility of f, relative to her sensitivity to her best-case expected utility of f, than is DM2.

---

[6]Ghirardato and Marinacci (2002) begin with this definition of comparative uncertainty aversion, and then go on to add an additional requirement to separate out risk aversion from ambiguity aversion. This is because they work within an Anscombe and Aumann (1963) framework, and is not important to our argument here.

3.2 Ambiguity attitude in CES meta-utility models

An immediate corollary of Proposition 2 is that the MRS coincides with the standard notion of ambiguity aversion for the α-MEU model. We state this as a general result for CES meta-utilities:

Corollary 3 *For a DM with a CES meta-utility function*

$$U(\boldsymbol{w}, \boldsymbol{b}) = [\alpha \boldsymbol{w}^\rho + (1 - \alpha) \boldsymbol{b}^\rho]^{1/\rho}$$

the DM is ambiguity averse (respectively, neutral, seeking) if and only if

$$\left(\frac{\alpha}{1 - \alpha} \right) \left(\frac{\boldsymbol{b}}{\boldsymbol{w}} \right)^{1-\rho} > \ \ (respectively, =, <) \ 1. \tag{1}$$

In particular, a DM with perfect substitutes (α-MEU, $\rho = 1$ and $\alpha < 1$) meta-utility is ambiguity averse if and only if $\alpha > 1/2$, and a DM with perfect complements (MEU, $\rho = -\infty$ or $\alpha = 1$) meta-utility is always ambiguity averse.

Proof.

Immediate from calculating the MRS and applying Proposition 2. ∎

We note here that ambiguity attitude in a CES meta-utility model depends on both ρ and α. Furthermore, unless $\rho = 1$ or $\rho \to -\infty$, the DM's ambiguity attitude is relative, in the sense that it depends on the ratio of **b** to **w**. As we discuss below in Subsection 3.3, these properties provide important distinctions and similarities between our model and ambiguity attitude as measured in the smooth model of Klibanoff, Marinacci and Mukerji (2005).

Figure 4 illustrates the relationship between the DM's attitude toward midpoint-preserving spreads and the MRS. Recall that the 45° **w** = **b** line forms the lower boundary of the domain of the meta-utility function. For any point (u, u) on this line, its midpoint-preserving spreads correspond to the points on the line segment from the y-axis to (u, u) with slope -1.

In the figure, the indifference curves depict a Binmore-Cobb-Douglas meta-utility, with $\alpha = 3/4$. Because $\alpha > 1/2$ (and $\mathbf{b} \geqslant \mathbf{w}$), expression (1) is always greater than one (note that $\rho = 0$ in the Binmore-Cobb-Douglas case). The highest indifference curve on a given midpoint-preserving spread line lies on the 45° line. This shows that the DM is globally ambiguity averse.

A convenient property of CES meta-utility functions is that they are continuously differentiable, so their MRS is well-defined. Definition 3 allows for cases in which the meta-utility function could have a kink. In this case, the MRS when approached from the left differs from the MRS when approached from the right. In this way, the intuition of Proposition 2 enables us to study a DM whose ambiguity attitude is not always sharp. Such a DM need not be viewed as irrational or inconsistent in our framework.
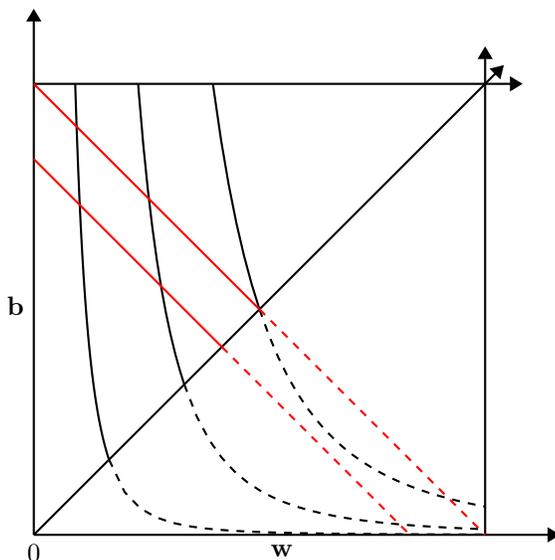
Figure 4: Binmore-Cobb-Douglas Indifference Curves with $\alpha = \frac{3}{4}$. Red lines reflect midpoint-preserving spreads.

3.3 Ambiguity attitude in the smooth and variational models

The models we have considered so far all depend on at most two sufficient statistics, **w** and **b**. We finish this section by discussing the relationships between our models and those that can depend on more information. In particular, we focus on the smooth model (Klibanoff, Marinacci and Mukerji, 2005) and the variational model (Maccheroni, Marinacci and Rustichini, 2006). We caution, however, that these models address different problems from ours. In their settings, the DM has a potentially large collection of priors, and calculates an expected utility with respect to each prior. The DM in these models has an abundance of information. We are instead interested in acts that are potentially non-measurable, for which the DM has very little information.

Nonetheless, we show that our approach can clarify aspects of ambiguity attitude as defined in these models. In particular, as we shall see, constant relative ambiguity attitude in the smooth model is closely related to our CES meta-utility models. Likewise, the variational model is, in an important but seemingly distinct sense, one of constant absolute ambiguity attitude (see the discussion in Grant and Polak, 2013), an idea that our meta-utility approach makes precise.

We begin by discussing the smooth model. In this model, the DM faces two sources of risk, with one interpreted as a probability distribution over priors and the other as a probability distribution given a prior. She has a first-order utility function u, which she uses to calculate her expected utility conditional on a prior, and a cardinal second-order utility function ϕ. We might therefore think of

13

the smooth model as applying the Savage axioms twice.

In their model, the DM is considered ambiguity averse (respectively neutral, seeking) if and only if the second-order utility ϕ is concave (respectively linear, convex). This provides a natural analogy between ambiguity attitude and risk attitude. On the other hand, the DM's beliefs about how likely a better (first-order) expected utility is, compared with a worse one, are unrelated to the DM's ambiguity attitude. In this way, a DM in the smooth model can be optimistic or pessimistic about how ambiguity will resolve, independently of whether the DM likes or dislikes ambiguity.

To illustrate, consider the following example:

Example 1 *Let f be an act that depends on a two-color Ellsberg urn, where the probability of drawing a red ball is p with second-order probability α and $q > p$ with probability $1 - \alpha$. Suppose that the act yields a higher prize when a red ball is drawn. Then, the DM receives her worst-case expected utility $\mathbf{w}$ with probability α and her best-case expected utility $\mathbf{b}$ with probability $1 - \alpha$. The smooth utility of f is*

$$V(f) = \alpha\phi(\mathbf{w}) + (1 - \alpha)\phi(\mathbf{b})$$

If we interpret this smooth utility as a function of $\mathbf{w}$ and $\mathbf{b}$, the marginal rate of substitution is

$$\left(\frac{\alpha}{1 - \alpha}\right)\left(\frac{\phi'(\mathbf{w})}{\phi'(\mathbf{b})}\right)$$

For example, if $\phi(v) = v^\rho$, then $\phi'(v) = \rho v^{\rho - 1}$, and the MRS becomes

$$\left(\frac{\alpha}{1 - \alpha}\right)\left(\frac{\mathbf{b}}{\mathbf{w}}\right)^{1 - \rho}$$

From Example 1, we see a clear analogy between the smooth model with Klibanoff, Marinacci and Mukerji's (2005) notion of a constant relative ambiguity attitude and our model's CES meta-utility. However, the interpretations are importantly distinct. In the smooth model, the parameter α is not part of the DM's ambiguity attitude; it represents her second-order prior belief. In the special case where $\rho = 1$, the DM is ambiguity neutral in the smooth model regardless of how pessimistic she is. A CES DM in our model with $\rho = 1$ is ambiguity neutral if and only if $\alpha = 1/2$. Which model is more appropriate depends on whether pessimism or optimism is an essential component of ambiguity attitude in a given context.

In the special case where $\alpha = 1/2$, the issue of pessimism or optimism is out of the picture. In this case, a smooth DM with globally strictly concave (respectively linear, strictly convex) ϕ always has $\phi'(\mathbf{w}) > \phi'(\mathbf{b})$ (respectively equal, less than). In that case, the smooth definition of ambiguity attitude and our MRS definition coincide.

Another way of putting the point is that our MRS definition of ambiguity attitude has two

components, whereas the smooth definition has only one. When one component (the pessimism / optimism) drops out—and when the smooth DM's ambiguity attitude is global—the two coincide in terms of whether the DM counts as ambiguity averse, neutral, or seeking.

The argument in Example 1 generalizes to ambiguous events with more than two expected utilities. We state this as follows:

Proposition 4 *Let f be a smooth-model ambiguous act, with possible expected utilities*

$$\mathbf{w} := u_1 < u_2 < \ldots < u_n := \mathbf{b}$$

and associated second-order probabilities

$$p_1, p_2, \ldots, p_n,$$

and suppose $p_1 = p_n$.

Define the associated meta-utility function

$$U(\mathbf{w}, \mathbf{b}) = p_1 \phi(\mathbf{w}) + p_n \phi(\mathbf{b})$$

Then if the DM is globally smooth-ambiguity averse (respectively neutral, seeking), the MRS is greater than (respectively equal to, less than) 1.

Proof. If $p_1 = p_n$, then the MRS of the meta-utility function is

$$\frac{\phi'(\mathbf{w})}{\phi'(\mathbf{b})}$$

If the DM is globally smooth-ambiguity averse, ϕ is strictly concave, so the MRS is greater than 1. The other cases are similar. ∎

We now turn to the variational model of Maccheroni, Marinacci and Rustichini (2006). In their setting, the DM evaluates an act $f \in \mathcal{F}$ with a set C of priors by balancing the expectation of a utility index u, given a prior, against a cost c associated with that prior:

$$V(f) = \min_{p \in C} \left\{ \int u(f) dp + c(p) \right\}$$

As in the MEU model, the DM is pessimistic. However, the DM also may view some possible priors as highly implausible. The variational model tempers the MEU idea of evaluating an ambiguous act pessimistically by essentially discounting less plausible priors.

For us, the crucial feature of the variational model is the way it captures ambiguity attitude;

it corresponds to quasi-linear preferences in our setting. We show this next, starting with the observation that variational DMs exhibit constant absolute ambiguity aversion.

The weak certainty independence axiom of the variational model implies the following condition, as Grant and Polak (2013) observe: for all weights $t \in (0,1)$, for any certain prizes $x, y, z \in X$, and for any act $f \in \mathcal{F}$,

$$\text{If } tf + (1-t)x \succsim tz + (1-t)x, \text{ then } tf + (1-t)y \succsim tz + (1-t)y$$

Thus, if the DM prefers a mixture of f with some constant prize x to the same mixture of prize z with x, then replacing x with y preserves the preference ordering. We can add or subtract any constant to the comparison of f with z, and the ordering is maintained. In this sense, the ambiguity attitude is independent of wealth effects.

To show that our model clarifies this sense of constant absolute ambiguity attitude, we consider the general class of mean-dispersion preferences, of which the variational model is an important example.[7] As Grant and Polak (2013) show, variational preferences can be expressed as

$$V(f) = \mu - \rho(d) \tag{2}$$

where μ is the average expected utility over all priors, d is utility dispersion, and $\rho(\cdot)$ is a cost function associated with utility dispersion.

This model is comparable to ours if $\mu = (\mathbf{b} + \mathbf{w})/2$ (the DM symmetrically averages over priors) and if $d = (\mathbf{b} - \mathbf{w})/2$ (utility dispersion is symmetric). In this case, the meta-utility is quasi-linear in $\mathbf{b}$ and $\mathbf{w}$:

$$(\forall k \in [-\mathbf{w}, 1 - \mathbf{b}]) \, U(\mathbf{w} + k, \mathbf{b} + k) = U(\mathbf{w}, \mathbf{b}) + k.$$

It is clear that the MRS is independent of any additive constant k. In this way, our model clarifies that the variational (and related) models indeed reflect constant absolute ambiguity aversion.

In sum, the MRS embeds important features of ambiguity attitude in both the smooth and variational models. Moreover, we gain new insight into these models and into what relative and absolute ambiguity attitudes mean.

4 Ambiguity with Strategic Interaction

We now show that the meta-utility approach provides novel insights in multi-player settings. In this section, we demonstrate that individuals who are globally ambiguity neutral still need to consider

---

[7]Several other models, such as Siniscalchi (2009), also have mean-dispersion representations. See Grant and Polak (2013).

ambiguity when interacting strategically. This is because updating beliefs under ambiguity affects both outer and inner measures, often in complex ways.[8]

We illustrate the distinction between ambiguity attitude and ambiguity sensitivity in a two-player game with debt contracting, in which a borrower (he) has private information about his endowment and can fraudulently claim to be unable to repay the full face value of a debt. The lender (she) can, at a cost, audit the borrower to verify whether a default is due to bad luck or fraud. If caught lying, the borrower owes the debt plus a penalty. This problem of costly verification, raised by Townsend (1979) and Gale and Hellwig (1985), has found applications in macroeconomics (Bernanke, Gertler and Gilchrist, 1996) and in allocation problems (Ben-Porath, Dekel and Lipman, 2014; Mylovanov and Zapechelnyuk, 2017). In the absence of a commitment device, the unique equilibrium is in mixed strategies. The lender mixes between auditing and trusting, and the borrower mixes between honesty and lying (see Reinganum and Wilde, 1986; Krasa and Villamil, 2000). If, as turns out to be the case, a greater probability of lying leads the lender to increase her audit probability, then all else equal, reducing the rate of strategic default improves efficiency.

The lender's equilibrium auditing probability depends on her posterior beliefs, conditional on default, that the borrower was a victim of bad luck and not a fraud. We now show that if the lender views this probability as ambiguous, the borrower's equilibrium lying probability may be lower than it would be under risk.

For concreteness, we restrict attention to a world with two payoff-relevant events, $\{E, \Omega \backslash E\}$. Call the lender A and the borrower B. In event E, B can afford to repay the debt in full; in $\Omega \backslash E$, he can give A only some smaller amount. Upon learning his endowment, B decides whether to default, either because E has not occurred or because B chooses to attempt a fraud.

If B does not default, he repays the debt in full to A, and the game ends. Otherwise, A decides whether to accept the reduced payment from B or to audit. If A audits and uncovers a fraud, B owes a penalty, which more than compensates for the audit cost. If A performs an audit and finds that B had defaulted due to bad luck, she gets what B had offered and loses the cost she spent on her audit. The game is depicted in Figure 5.

When deciding whether to default, B faces no uncertainty about his endowment. Accordingly, we restrict our focus to A's payoffs and their effects on B's probability p of lying. The best outcome for A is for her to uncover a fraud and receive her payment plus the penalty; normalize her utility of this outcome to 1. The worst outcome for A is for her to perform an audit and come up empty; normalize the utility of this outcome to 0. Finally, if she trusts B and accepts his reduced payment, A receives utility $\hat{u} \in (0, 1)$.

As a benchmark, consider a setting of risk with a common prior $\pi = Pr(E)$ that B can fully
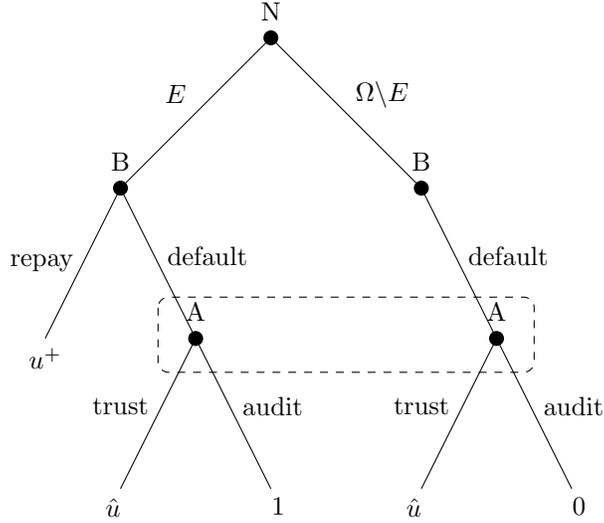

Figure 5: Game tree for costly state verification problem. Only A's payoffs are shown; they are such that $0 < \hat{u} < u^+ < 1$.

repay the debt. Conditional on default and on B's lying probability p, A's expected utility from auditing is

$$W_A\left(\text{audit}\right) = \frac{\pi p}{1 - \pi + \pi p} \tag{3}$$

Her expected utility of trusting B is $\hat{u}$. Proposition 5 gives B's equilibrium lying probability p^*.

Proposition 5 *Under risk, if there is an interior solution, then the borrower lies with probability*

$$p^* = \left(\frac{1 - \pi}{\pi}\right)\left(\frac{\hat{u}}{1 - \hat{u}}\right) \tag{4}$$

If p^* in (4) is greater than or equal to one, then the market shuts down.

Consequently, as the probability π of E increases, the probabilities of both strategic and inherent default decrease.

Proof. Immediate from setting (3) equal to $\hat{u}$. ∎

Next, consider a setting with ambiguity. Let $\underline{\pi}$ be A's inner measure of event E, and let $\bar{\pi}$ be her outer measure of E. Assume $0 < \underline{\pi} < \bar{\pi} < 1$. In light of Proposition 5, we know that A's best-case expected utility $\mathbf{b}$ is found at $\bar{\pi}$ and her worst-case expected utility $\mathbf{w}$ is found at $\underline{\pi}$.

Suppose A has CES meta-utility and is globally ambiguity neutral:

$$U_A(\mathbf{w}, \mathbf{b}) = \frac{\mathbf{w} + \mathbf{b}}{2}$$

We now show that B's equilibrium lying probability can be lower than (4). Thus, ambiguity can

make lending more attractive, a result that seems unknown in the literature on ambiguity and financial markets.

Substituting for $\mathbf{w}$ and $\mathbf{b}$, A's meta-utility of an audit given default and given p is

$$\left(\frac{\underline{\pi}p}{1 - \underline{\pi} + \underline{\pi}p} + \frac{\bar{\pi}p}{1 - \bar{\pi} + \bar{\pi}p} \right) / 2 \tag{5}$$

Setting (5) equal to $\hat{u}$ gives the following quadratic equation for the optimal lying probability p^{**}:

$$-2(1 - \hat{u})\underline{\pi}\bar{\pi}(p^{**})^2 + (2\hat{u} - 1)(\underline{\pi} + \bar{\pi} - 2\underline{\pi}\bar{\pi})p^{**} + 2\hat{u}(1 - \underline{\pi})(1 - \bar{\pi}) = 0$$

In the special case where $\hat{u} = 1/2$, p^{**} has the following simple form:

$$p^{**} = \left[\frac{(1 - \underline{\pi})(1 - \bar{\pi})}{\underline{\pi}\bar{\pi}} \right]^{1/2} \tag{6}$$

Proposition 6 gives the result that $p^{**} < p^*$, i.e., that there is less lying and hence greater efficiency under ambiguity than under risk:

Proposition 6 *Assume $\hat{u} = 1/2$. Let $\hat{\pi} = (\underline{\pi} + \bar{\pi})/2$. Let p^* be the lying probability under risk (4) given $\pi = \hat{\pi}$, and let p^{**} be the lying probability under ambiguity (6). Then:*

1. If $\hat{\pi} > 1/2$, then $p^{**} < p^*$.

2. If $\hat{\pi} \leqslant 1/2$, then the market shuts down.

Proof. Define $r = (\bar{\pi} - \underline{\pi})/2$. Then $\underline{\pi} = \hat{\pi} - r$ and $\bar{\pi} = \hat{\pi} + r$. Substituting into p^{**} gives

$$p^{**} = \left[\frac{(1 - \hat{\pi} + r)(1 - \hat{\pi} - r)}{(\hat{\pi} - r)(\hat{\pi} + r)} \right]^{1/2}$$
$$= \left[\frac{(1 - \hat{\pi})^2 - r^2}{\hat{\pi}^2 - r^2} \right]^{1/2} \tag{7}$$

If $\hat{\pi} \leqslant 1/2$, the numerator in (7) is easily seen to be at least as large as the denominator, and since the only equilibrium is in mixed strategies, the market shuts down and the lying probability is latent.

If $\hat{\pi} > 1/2$, then (7) is below one. In this case, we have $p^* > p^{**}$ if and only if

$$\frac{(1 - \hat{\pi})^2 - r^2}{\hat{\pi}^2 - r^2} > \frac{(1 - \hat{\pi})^2}{\hat{\pi}^2} \qquad \Leftrightarrow \qquad \hat{\pi} > \frac{1}{2}$$

Hence, if the market does not shut down (i.e., if $\hat{\pi} > 1/2$), then the lying probability is lower under ambiguity than the corresponding lying probability under risk. ∎

19

References

Agranov, Marina, and Andrew Schotter. 2012. "Ignorance Is Bliss: An Experimental Study of the Use of Ambiguity and Vagueness in the Coordination Games with Asymmetric Payoffs." *American Economic Journal: Microeconomics*, 4(2): 77–103.

Ahn, David, Syngjoo Choi, Douglas Gale, and Shachar Kariv. 2014. "Estimating Ambiguity Aversion in a Portfolio Choice Experiment." *Quantitative Economics*, 5(2): 195–223.

Anscombe, Francis J, and Robert J Aumann. 1963. "A definition of subjective probability." *The Annals of Mathematics and Statistics*, 34: 199–205.

Arrow, Kenneth J, and Leonid Hurwicz. 1972. "An Optimality Criterion for Decision-making under Ignorance." In *Uncertainty and Expectations in Economics.* Vol. 1, , ed. CF Carter and JL Ford. Oxford, England:Basil Blackwell & Mott Ltd.

Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman. 2014. "Optimal allocation with costly verification." *American Economic Review*, 104(12): 3779–3813.

Bernanke, Ben, Mark Gertler, and Simon Gilchrist. 1996. "The Financial Accelerator and the Flight to Quality." *Review of Economic Studies*, 78(1): 1–15.

Binmore, Kenneth G. 2009. *Rational Decisions.* Princeton, NJ:Princeton University Press.

Chateauneuf, Alain, Jürgen Eichberger, and Simon Grant. 2007. "Choice under uncertainty with the best and worst in mind: Neo-additive capacities." *Journal of Economic Theory*, 137(1): 538–567.

De Castro, Luciano, and Nicholas C Yannelis. 2018. "Uncertainty, efficiency and incentive compatibility: Ambiguity solves the conflict between efficiency and incentive compatibility." *Journal of Economic Theory*, 177: 678–707.

Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, 75(4): 643–669.

Frick, Mira, Ryota Iijima, and Yves Le Yaouanq. 2020. "Objective Rationality Foundations for (Dynamic) α-MEU." Cowles Foundation for Research in Economics Discussion paper 2244.

Gale, Douglas, and Martin Hellwig. 1985. "Incentive-Compatible Debt Contracts: The One-Period Problem." *Review of Economic Studies*, 52(4): 647–663.

Ghirardato, Paolo, and Marciano Siniscalchi. 2012. "Ambiguity in the Small and in the Large." *Econometrica*, 80(6): 2827–2847.

Ghirardato, Paolo, and Massimo Marinacci. 2002. "Ambiguity Made Precise: A Comparative Foundation." *Journal of Economic Theory*, 102: 251–289.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci. 2004. "Differentiating Ambiguity and Ambiguity Attitude." *Journal of Economic Theory*, 118(2): 133–173.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics*, 18: 141–153.

Good, Irving John. 1966. "Symposium on Current Views of Subjective Probability: Subjective Probability As the Measure of a Non-Measurable Set." In *Studies in Logic and the Foundations of Mathematics*. Vol. 44, 319–329. Elsevier.

Good, Irving John. 1983. *Good Thinking: The Foundations of Probability and Its Applications.* University of Minnesota Press.

Grant, Simon, and Ben Polak. 2013. "Mean-dispersion Preferences and Constant Absolute Uncertainty Aversion." *Journal of Economic Theory*, 148(4): 1361–1398.

Grant, Simon, Patricia Rich, and Jack Stecher. 2020. "Bayes and Hurwicz without Bernoulli." *Journal of Economic Theory*, Forthcoming.

Grant, Simon, Patricia Rich, and Jack Stecher. 2021. "Objective and Subjective Rationality and Decisions with the Best and Worst Case in Mind." *Theory and Decision*, Forthcoming.

Gul, Faruk, and Wolfgang Pesendorfer. 2014. "Expected Uncertainty Utility." *Econometrica*, 82(1): 1–39.

Gul, Faruk, and Wolfgang Pesendorfer. 2015. "Hurwicz expected utility and subjective sources." *Journal of Economic Theory*, 159: 465–488.

Guo, Huiyi, and Nicholas C Yannelis. 2021. "Full implementation under ambiguity." *American Economic Journal: Microeconomics*, 13(1): 148–178.

Hurwicz, Leonid. 1951. "Optimality Criteria for Decision Making Under Ignorance." Cowles Commission Discussion Paper, Statistics 370.

Izhakian, Yehuda. 2017. "Expected Utility with Uncertain Probabilities Theory." *Journal of Mathematical Economics*, 69: 91–103.

Jaffray, Jean-Yves, and Peter P Wakker. 1993. "Decision Making with Belief Functions: Compatibility and Incompatibility with the Sure-Thing Principle." *Journal of Risk and Uncertainty*, 7(3): 255–271.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A smooth model of decision making under ambiguity." *Econometrica*, 73(6): 1849–1892.

Kopylov, Igor. 2009. "Choice Deferral and Ambiguity Aversion." *Theoretical Economics*, 4(2): 199–225.

Krasa, Stefan, and Anne P Villamil. 2000. "Optimal Contracts When Enforcement is a Decision Variable." *Econometrica*, 68(1): 119–134.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica*, 74(6): 1447–1498.

Machina, Mark J. 2014. "Ambiguity Aversion with Three or More Outcomes." *American Economic Review*, 104(12): 3814–3840.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory.* Oxford University Press.

Mylovanov, Tymofiy, and Andriy Zapechelnyuk. 2017. "Optimal allocation with ex post verification and limited penalties." *American Economic Review*, 107(9): 2666–94.

Pintér, Miklós. 2019. "Objective ambiguity." Budapest University of Technology and Economics Working paper.

Reinganum, Jennifer F, and Louis L Wilde. 1986. "Equilibrium Verification and Reporting Policies in a Model of Tax Compliance." *International Economic Review*, 27(3): 739–760.

Rigotti, Luca, Chris Shannon, and Tomasz Strzalecki. 2008. "Subjective Beliefs and ex ante Trade." *Econometrica*, 76(5): 1167–1190.

Rothschild, Michael, and Joseph E Stiglitz. 1970. "Increasing Risk: I. A Definition." *Journal of Economic Theory*, 2(3): 225–243.

Savage, Leonard J. 1954. *The Foundations of Statistics.* John Wiley & Sons.

Seo, Kyoungwon. 2009. "Ambiguity and Second-Order Belief." *Econometrica*, 77(5): 1575–1605.

Siniscalchi, Marciano. 2009. "Vector Expected Utility and Attitudes Toward Variation." *Econometrica*, 77(3): 801–855.

Suppes, Patrick. 1974. "The Measurement of Belief." *Journal of the Royal Statistical Society, Series B (Methodological)*, 36(2): 160–191.

Townsend, Robert M. 1979. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory*, 21(2): 265–293.

Vierø, Marie-Louise. 2012. "Contracting in Vague Environments." *American Economic Journal: Microeconomics*, 4(2): 104–130.

Wald, Abraham. 1950. *Statistical Decision Functions.* New York:Wiley.