# Decision Under Ambiguity via Intermediate Microeconomics


Simon Grant
Australian National University, simon.grant@anu.edu

Patricia Rich
University of Bayreuth, patricia.rich@uni-bayreuth.de

Jack Stecher
University of Alberta, stecher@ualberta.ca



The goal of this paper is to present a unifying theory of decision-making under ambiguity. We present a theory that nests several of the existing models and captures special cases of others. Our theory requires no more technical background than found in intermediate microeconomics, and provides a natural interpretation of ambiguity attitude (it corresponds to a marginal rate of substitution). We illustrate the insights our model generates by applying it to a problem of lending with possible strategic default, in which the borrower's future wealth is ambiguous. To show that our theory rests on a solid foundation, we then provide an axiomatization. To summarize, our model generalizes existing theories, has a low barrier to entry, and is straightforward to work with.



**JEL Classification:** D80, D81

*Key words*: ambiguity, expected utility, optimism, pessimism, costly state verification


## 1. Introduction

The theory of decision under ambiguity, in contrast to that of decision under risk, lacks a general, unified framework. A synopsis of the sprawling landscape is given in Etner et al. (2012, 234): "The non-specialist observer might be overwhelmed by the number of different models aimed at capturing how ambiguity can affect decisions. And even before that, she might be baffled by terminology issues." The diversity of models has consequences for our ability to apply the theory of ambiguity to practical problems, since different models can support different conclusions.

A general approach would help clarify which issues are model-dependent, and what insights generalize (and to what extent, in which ways, and so forth). Ideally, we would also have an approach that is both easy to grasp and to work with—also for those outside of the theoretical decision theory community—and rich enough to accommodate the heterogeneous ambiguity attitudes that we know real-world agents have (see Etner et al. 2012, for an overview).

In what follows, we employ two familiar ingredients from intermediate microeconomics—expected utility and a two-good ordinal utility function which we call a *meta-utility function*—to provide a theory that unifies several existing models and sheds light on others. We then axiomatize this theory. Although the context is new, and the axiomatization necessarily includes some technicalities, our aim is to reassure you that decision under ambiguity is less exotic than it might seem. Or to express it in culinary terms, ambiguity tastes like chicken.

We use expected utility first, for those settings in which a decision maker (whom we designate as the DM throughout) has sufficient understanding of the underlying random process to quantify precisely the uncertainty she faces. In situations of ambiguity, however, as in Ellsberg's (1961) classic example, unique probabilities are not available. Nevertheless, the DM can bound the probability of any event. Accordingly, we can characterize each ambiguous act with *two* expected utilities: the worst-case and the best-case expected utility. Label these $\mathbf{w}$ and $\mathbf{b}$, respectively.[1]

These expected utilities behave like goods for our two-good ordinal utility function. In intermediate microeconomics, the ordinal utility might take wine and beer as inputs, and express how a DM values the trade-offs between them. In the theory we present below, the ordinal utility operates at a meta-level, taking $\mathbf{w}$ and $\mathbf{b}$ as the inputs. We call the resulting utility function an *ordinal Hurwicz expected utility* (OHEU). The term refers to Hurwicz's (1951) model, in which the DM calculates a linearly weighted average of $\mathbf{w}$ and $\mathbf{b}$. Our theory extends this by finding a general, not necessarily linear, way to combine $\mathbf{w}$ and $\mathbf{b}$.[2]

The meta-utility approach 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 (Wald 1950, Gilboa and Schmeidler 1989) views a DM as pessimistic, maximizing the worst-case expected utility representation, so that $U(\mathbf{w}, \mathbf{b}) = \min(\mathbf{w}, \mathbf{b}) =$

[1] That is, we adopt the measure-theoretic perspective on ambiguity, using outer and inner probabilities when the two do not converge. This perspective is pursued in numerous theories of decision under ambiguity; examples include Good (1966, 1983), Suppes (1974), Manski (1981), Binmore (2009), and Pintér (2022). In the extreme case in which the inner probability of an event is zero and the outer probability is one, the DM is completely ignorant in the sense of Arrow and Hurwicz (1972) (see also Schlag and Zapechelnyuk 2020). 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 et al. 2007, Grant and Polak 2013).

[2] Gul and Pesendorfer (2014) go in the opposite direction, applying what we interpret as a cardinal *interval* utility state-by-state to a worst and best outcome, then aggregating. In the special case of linear averaging, as in Hurwicz (1951), their approach coincides with ours.

**Figure 1 Leontief Meta-utilities (Gilboa-Schmeidler).**

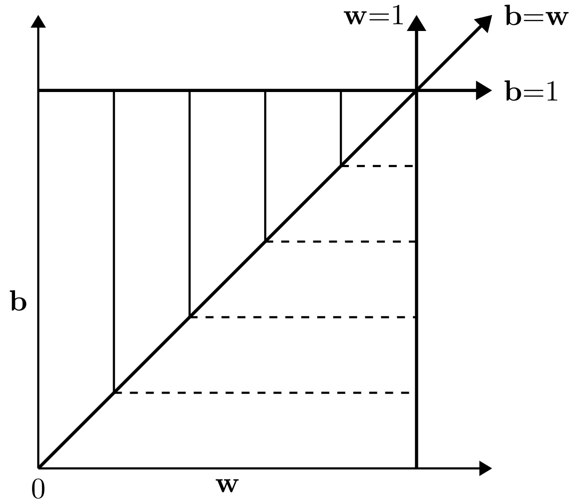


*Note.* $U(\mathbf{w},\mathbf{b}) = \min\{\mathbf{w},\mathbf{b}\} = \mathbf{w}$. In the general CES equation, this corresponds to $\alpha = 1$ and $\rho = 1$ or $-\infty$.

$\mathbf{w}$. The $\alpha$-maxmin expected utility ($\alpha$-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 a constant rate of substitution $\alpha \in [0,1]$ on $\mathbf{w}$ reflecting the DM's degree of pessimism. The geometric $\alpha$-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 worst and best outcome, we can normalize their utilities to 0 and 1 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.

These CES models of ambiguity are undoubtedly simpler than some other models in the literature (such as the smooth and variational models, discussed later). They apply to cases in which a DM has very little probabilistic information; models such as smooth and variational, on the other hand, apply to settings in which the DM has considerable information. Real-world choice problems in which agents have too little information to measure the probabilities of some important events seem to be very common. Examples include cases for which prediction is known to be difficult

**Figure 2 Perfect Substitutes Meta-utilities (Hurwicz).**

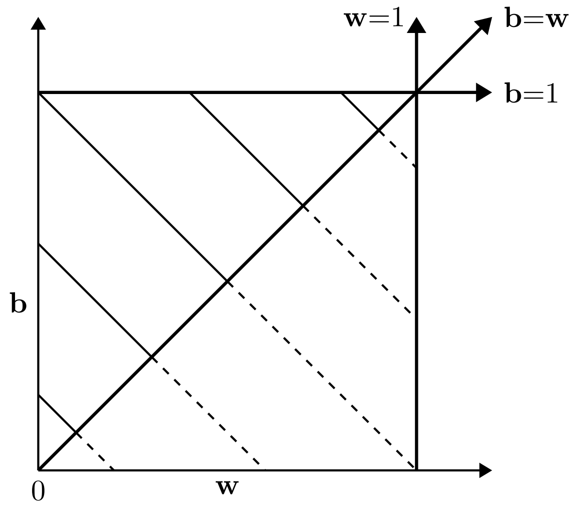


*Note.* $U(\mathbf{w},\mathbf{b}) = \alpha\mathbf{w} + (1-\alpha)\mathbf{b}$. In the general CES equation, Hurwicz preferences correspond to $\alpha < 1$ and $\rho = 1$. The figure shows the case of $\alpha = \frac{1}{2}$.

**Figure 3 Cobb-Douglas Meta-utilities (Binmore).**

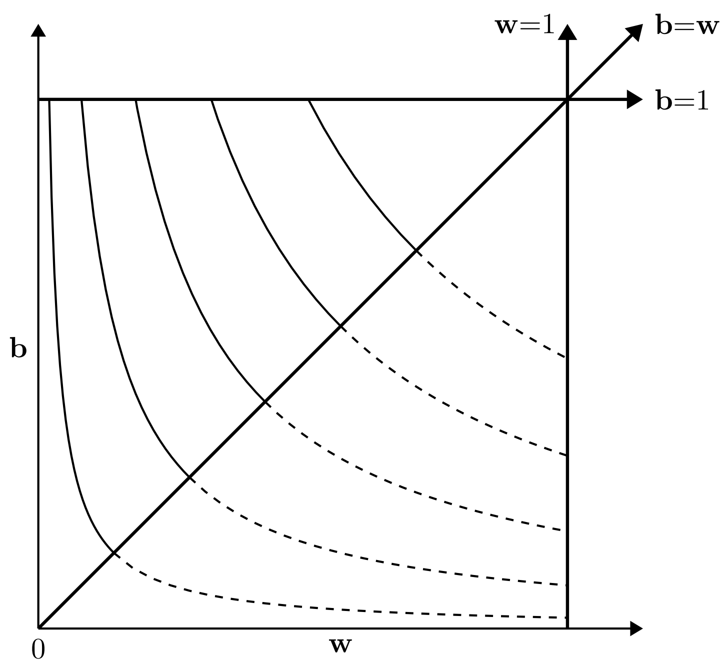


*Note.* $U(\mathbf{w},\mathbf{b}) = \mathbf{w}^{\alpha} + \mathbf{b}^{1-\alpha}$. In the general CES equation, these preferences correspond to $\rho = 0$. The figure shows the case of $\alpha = \frac{1}{2}$.

even for experts, for example if the event in question is an election result or other social/political outcome in the non-near future. Real agents may also treat such events as being non-measurable even if we think an ideal agent could in principle do better. Furthermore, even a DM who entertains multiple models that could capture the relevant causal processes (as Berger et al. 2017, discuss for the case of climate change) may not be able to place sensible weights on those different models; in such cases, it is very natural for the DM to focus on a worst case and to temper it with a best case (or an average case; see Section 2). Hence, we see the choice of model as something that should depend on the intended application, and simplicity is often a virtue.

A further advantage of the simplicity of our model is that it can help us to clarify the concept of ambiguity attitude. In our OHEU approach, ambiguity attitude has a natural characterization that comes directly from the meta-utility function. The marginal rate of substitution (MRS) between $\mathbf{w}$ and $\mathbf{b}$ captures how the DM views the trade-off between the security level $\mathbf{w}$ and the aspiration level $\mathbf{b}$. This not only extends the utility-based approach that is common in $\alpha$-MEU models, but also, as we show, captures essential aspects of ambiguity attitude as defined in models based on a probabilistic approach.

The concept of ambiguity attitude is in need of clarification because the literature on ambiguity makes use of a variety of notions, including relative and absolute ambiguity attitude, endogenous and exogenous definitions, and otherwise different conceptions resulting from different theories of how individuals respond to ambiguity. For example, in Bossaerts et al. (2010) and Asparouhova et al. (2015), ambiguity neutrality in an $\alpha$-MEU setting is defined as taking $\alpha = 1/2$, reflecting neither optimism nor pessimism. Ghirardato et al. (2004) treats ambiguity neutrality as expected utility maximization, acting as if there is a unique prior. Further complicating these distinct notions of ambiguity attitude is that the former is in utility space, and the latter is instead about perceived probabilities. Ahn et al.'s (2014) experiments treat ambiguity attitude in a model-dependent way, defined in terms of the pessimism parameter $\alpha$ in the $\alpha$-MEU model and in a somewhat more complex way in the smooth model (Klibanoff et al. 2005). A way of systematizing and comparing these various approaches would be very helpful. When viewed in terms of the MRS of the meta-utility function, many of the similarities in the different notions of ambiguity attitude become more readily apparent, and where these approaches differ, it becomes easier to focus on their distinctions. Thus, the approach provides an over-arching perspective on the literature which was previously lacking.

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 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 that ambiguity is irrelevant. These effects would not be possible to capture under a notion of ambiguity neutrality as expected utility preferences.

The rest of this article is organized as follows: Section 2 discusses alternative notions of ambiguity attitude and illustrates how the MRS can be viewed as a general definition of (local) ambiguity attitude. Section 3 shows a concrete application of our approach to a strategic setting with costly state verification. Section 4 shows that the meta-utility theory is well-founded, specifying a formal model, presenting the axioms, and stating a representation result. Section 5 briefly concludes.

## 2. A General Definition of Ambiguity Attitude: The MRS

### 2.1. Alternative notions of ambiguity attitude

In order to further motivate our own characterization of ambiguity attitude, it will be helpful to elaborate on the various alternatives. There are many intuitive ways of thinking about ambiguity attitude; different modeling approaches reflect different underlying intuitions. Ellsberg (1961, 661–664) describes multiple ways that one might evaluate the options in his 3-color urn problem—presaging the suite of existing modern models—on the way to supplying his own proposal (Ellsberg 1961, 665–668). Ellsberg's proposed decision criterion involves taking a linearly-weighted average of a worst-case expected utility $\mathbf{w}$ and an estimated expected utility that the DM deems most likely.[3] This is equivalent to the $\alpha$-MEU model, with a modified weight. Ellsberg uses the weight on the worst-case expected utility as a measure of conservatism, and states that as this weight vanishes, the DM's most likely estimate dominates the decision.

We note that Ellsberg's criterion presupposes a set $Y_0$ of priors that the DM considers reasonable, and holds this set fixed. Thus, while it is natural to consider a DM who simply maximizes expected utility to be ambiguity neutral, in that they ignore all but one most subjectively plausible prior, in another sense this DM still faces ambiguity. The weight on the worst-case prior shrinks to zero, but the DM still considers the entire set $Y_0$ as reasonable. Nevertheless, a DM who is indistinguishable from an expected utility maximizer is observationally equivalent to one for whom $Y_0$ contains a single distribution. This definition of ambiguity neutrality is not the only natural one, but it is the most intuitive one to use in several models. This includes the smooth model (Klibanoff et al. 2005), which essentially involves applying expected utility twice (once for each possible prior, and then again using a cardinal second-order utility to aggregate over the resultant first-order expected utilities). We look at the smooth model in more detail below.

[3] In his Nobel Lecture, Hansen (2014, 977) uses this approach as a tractable way to account for model uncertainty, in addition to discussing variational preference (Hansen and Sargent 2001, Maccheroni et al. 2006) and the smooth model (Klibanoff et al. 2005).

In other models, such as $\alpha$-MEU, Ellsberg's notion of conservatism—the amount of weight put on a worst-case expected utility—is a more natural measure of ambiguity attitude. Given Ellsberg's picture, a DM who puts weight on a worst-case expected utility relative to a neutral benchmark is naturally labeled as ambiguity averse, while one who focuses on a neutral benchmark is naturally labeled as ambiguity neutral. The measure-theoretic perspective gives the DM only a lower and upper expected utility, so in an $\alpha$-MEU setting, ambiguity neutrality would correspond to a DM with $\alpha = 1/2$. This is the viewpoint that Bossaerts et al. (2010) and Asparouhova et al. (2015) take.

A striking difference between these conceptions of ambiguity attitude is revealed by an example raised in Becker and Brownson (1964) and Einhorn and Hogarth (1986), attributed to Ellsberg.

Example 1. Consider a risky act $g$ that depends on an urn containing 1000 balls, labeled sequentially from 1 to 1000. The act $g$ gives prize $m$ if a specific ball, say number 382, is drawn. Otherwise, $g$ gives prize $\ell$. If all balls are equally likely to be drawn and we normalize $v(\ell) = 0$ and $v(m) = 1$, then the expected utility of $g$ is $0.999v(\ell) + 0.001v(m) = 0.001$.

In Example 1, compare $g$ with an ambiguous act $f$, which pays $m$ if 382 is drawn from a different urn and $\ell$ otherwise (we can imagine the number 382 to have been chosen by the DM to avoid concerns about foul play, as in Kadane 1992). The urn on which $f$ depends also contains 1000 balls, each with a number in $\{1, \ldots, 1000\}$, but has an unknown composition. Given the normalization of $v(\ell) = 0$ and $v(m) = 1$, a DM with $\alpha$-MEU preference representation strictly prefers $f$ to $g$ as long as $\alpha < 0.999$. On the other hand, a DM in the smooth model might reasonably believe that any composition of the ambiguous urn is as likely as any other. In that case, if the second-order utility is strictly concave, the smooth DM would always prefer $g$ to $f$.

These two perspectives illustrate distinct aspects of ambiguity attitude. The Hurwicz $\alpha$-MEU DM can be ambiguity averse ($\alpha > 1/2$), while strictly preferring an ambiguous urn to what we can loosely describe as an analogous risky urn. Act $f$ has a bound on its expected utility that is extremely close to the expected utility of act $g$, but an upper bound that is much larger. The Hurwicz perspective is that only an extremely ambiguity averse DM would strictly prefer $g$ to $f$, because the potential gain from choosing $f$ over $g$ is so much larger than the potential loss. On the other hand, a smooth DM could not be globally ambiguity averse and prefer an ambiguous urn to an analogous risky urn.

It would be useful to have a definition of ambiguity attitude that generalizes both these perspectives. To do so, we work in utility space. Thus, for the smooth model, although ambiguity neutrality collapses to expected utility, it is important to state the definition in terms of an affine second-order utility function, so that we can compare with models such as $\alpha$-MEU. Beyond the theoretical value of a unified viewpoint on ambiguity attitude, such a perspective can be seen to be important in applications of the theory of ambiguity. This includes both microeconomic applications and experimentation to understand individual choice behavior, and macroeconomic applications (as discussed in Section 3).

Thus, in what follows, we show how to use the meta-utility model as a general framework for studying ambiguity attitude, making it possible to separate out distinct aspects and components thereof. The proposed general definition of ambiguity attitude will be shown to capture the intuitions underlying many other proposals, both those which the meta-utility was obviously designed to generalize (such as Hurwicz' theory) and those which on the surface appear very different.

### 2.2. 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:

DEFINITION 1. Let $f, g \in \mathcal{F}$ be two acts with corresponding expected utility bounds $(\mathbf{w}^f, \mathbf{b}^f)$ and $(\mathbf{w}^g, \mathbf{b}^g)$. Say that $f$ is a *midpoint-preserving spread* of $g$ if, for some $\delta > 0$,

$$\mathbf{w}^g - \mathbf{w}^f = \delta = \mathbf{b}^f - \mathbf{b}^g.$$

If $f$ is a midpoint-preserving spread of $g$, then we view $f$ as more ambiguous than $g$.[4]

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.[5]

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

[4] In the notation used in Section 4, where the full model is presented, the utility bounds would be expressed as $(\underline{v}_\pi^f, \bar{v}_\pi^f)$ and $(\underline{v}_\pi^g, \bar{v}_\pi^g)$ and the definition of a midpoint-preserving spread would be $\underline{v}_\pi^g - \underline{v}_\pi^f = \delta = \bar{v}_\pi^f - \bar{v}_\pi^g$.

[5] A related notion of focusing on midpoints is in Baillon et al. (2012), which Abdellaoui and Zank (2023) apply to models with source-dependent preferences. 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 et al. (2004), Rigotti et al. (2008), and Ghirardato and Siniscalchi (2012).

PROPOSITION 1. *Let $f \in \mathcal{F}$ be an act with worst-case expected utility $\boldsymbol{w}$ and best-case expected utility $\boldsymbol{b}$. 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 of Proposition 1* 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,

$$\begin{aligned}\frac{\partial U}{\partial r} \lesseqqgtr 0 &\iff \frac{\partial U}{\partial \mathbf{b}} \lesseqqgtr \frac{\partial U}{\partial \mathbf{w}}\\ &\iff \frac{\partial U/\partial \mathbf{w}}{\partial U/\partial \mathbf{b}} \gtreqqless 1\end{aligned}$$

□

Proposition 1 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 $\succsim_2$.

Thanks to Proposition 1, 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.

### 2.3. Ambiguity attitude in CES meta-utility models

An immediate corollary of Proposition 1 is that the MRS coincides with the standard notion of ambiguity aversion for the $\alpha$-MEU model. We state this as a general result for CES meta-utilities:

COROLLARY 1. *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*

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

*In particular, a DM with perfect substitutes ($\alpha$-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 of Corollary 1* Immediate from calculating the MRS and applying Proposition 1. □

We note here that ambiguity attitude in a CES meta-utility model depends on both $\rho$ and $\alpha$. 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 2.4, these properties provide important distinctions and similarities between our model and ambiguity attitude as measured in the smooth model of Klibanoff et al. (2005).

Figure 4 illustrates the relationship between the DM's attitude toward midpoint-preserving spreads and the MRS. Recall that the 45° $\mathbf{w} = \mathbf{b}$ 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.[6]

[6] A convenient property of CES meta-utility functions is that they are continuously differentiable, so their MRS is well-defined. Nothing, however, precludes the indifference map generated by the meta-utility aggregator exhibiting kinks (especially along the 45 degree line, points for which **w** and **b** are equal). In such a case, the MRS when approached from the left differs from the MRS when approached from the right. In this way, the intuition of Proposition 1 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 $\alpha = \frac{3}{4}$.**

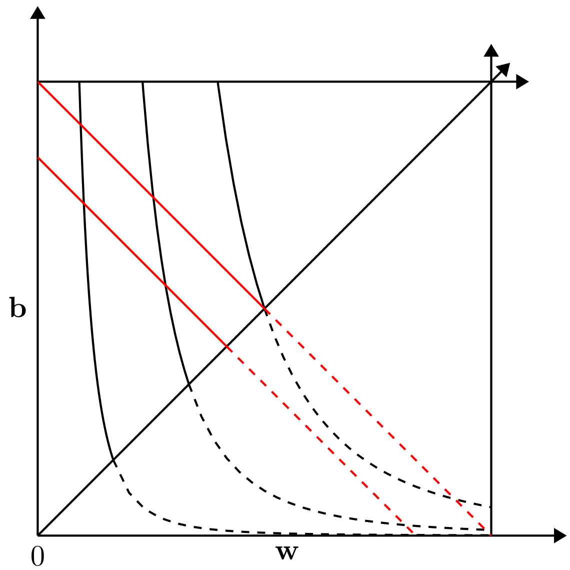


*Note.* Red lines reflect midpoint-preserving spreads.

## 2.4. Ambiguity attitude in the smooth and variational models

The models we have considered so far all depend on at most two sufficient statistics, $\mathbf{w}$ and $\mathbf{b}$. 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 et al. 2005) and the variational model (Maccheroni et al. 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 $\phi$. We might therefore think of 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 $\phi$ 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 (see again Example 1 in Subsection 2.1). 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 how we can understand the smooth model through the lens of the MRS, consider the following example:

EXAMPLE 2. 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 $\alpha$ 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 $\alpha$ 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 2, we see a clear analogy between the smooth model with Klibanoff et al.'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 $\alpha$ 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) $\phi$ always has $\phi'(\mathbf{w}) > \phi'(\mathbf{b})$ (respectively $=$, $<$). In that case, the smooth definition of ambiguity attitude and our MRS definition coincide.

Thinking back to the CES meta-utility case above, our MRS definition of ambiguity attitude has two components: The first, $\alpha$, represents pessimism (or ambiguity attitude as in Hurwicz, 1951, and Gul and Pesendorfer, 2015, and as alluded to in Ellsberg, 1961). The second, $\rho$, captures smooth ambiguity attitude as in the smooth model of Klibanoff et al. (2005), and the curvature of the indifference curves in our representations. When one parameter, $\alpha$, 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 difference in conceptions of ambiguity attitude across models is characterized by which of these parameters is in focus.

The argument in Example 2 generalizes to ambiguous events with more than two expected utilities. We state this as follows:

Proposition 2. *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 of Proposition 2* 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, $\phi$ is strictly concave, so the MRS is greater than 1. The other cases are similar. □

We now turn to the variational model of Maccheroni et al. (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.[7]

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 $f \in \mathcal{F}$,

$$tf + (1-t)x \succsim tz + (1-t)x \qquad \Longrightarrow \qquad tf + (1-t)y \succsim tz + (1-t)y$$

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.[8] As Grant and Polak (2013) show, variational preferences can be expressed as

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

where $\pi$ 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 $\pi = (\mathbf{b}+\mathbf{w})/2$ (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.

[7] A similar idea is in Wang (2022), who proposes a maxmin RDU approach rather than maxmin EU.

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

## 3. Application: 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 ambiguity when interacting strategically. This is because updating beliefs under ambiguity affects both outer and inner measures, often in complex ways.[9]

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 et al. 1996) and in allocation problems (Ben-Porath et al. 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

[9] For example, in Binmore (2009), updating under ambiguity can introduce new symmetric equilibria into Battle-of-the-Sexes games that strictly Pareto dominate the only symmetric equilibrium under risk. In Vierø (2012), differences in pessimism in an $\alpha$-MEU model introduce a new variable for contracting. As a final example, De Castro and Yannelis (2018) show that bilateral bargaining under ambiguity in an MEU setting can achieve first-best outcomes.

**Figure 5 Game tree for costly state verification problem**

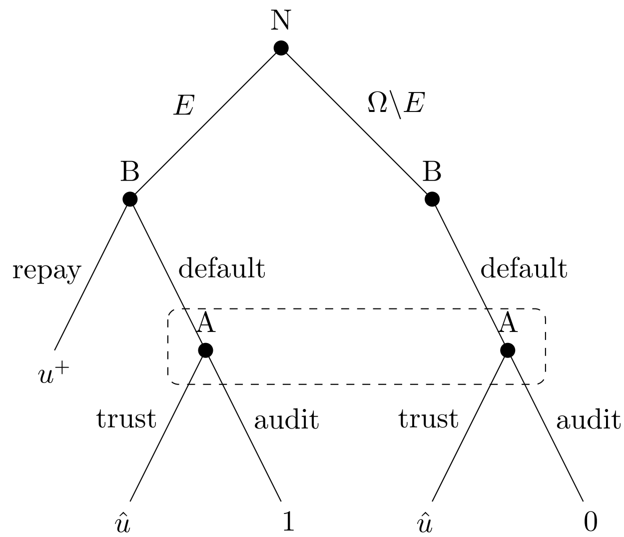


*Note.* Only A's payoffs are shown; they are such that $0 < \hat{u} < u^+ < 1$

for $A$ is for her to uncover a fraud and receive her payment plus the penalty; normalize her utility of this outcome to 1. As the penalty more than covers $A$'s audit cost, this is greater than the utility $u^+$ she would have received if $A$ had repaid her in full from the start. 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 repay the debt. Conditional on default and on $B$'s lying probability $p$, $A$'s expected utility from auditing is

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

Her expected utility of trusting $B$ is $\hat{u}$. Proposition 3 gives $B$'s equilibrium lying probability $p^*$.

Proposition 3. *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 $\pi$ of $E$ increases, the probabilities of both strategic and inherent default decrease.*

*Proof of Proposition 3* 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 3, 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 4 gives the result that $p^{**} < p^*$, i.e., that there is less lying and hence greater efficiency under ambiguity than under risk:

Proposition 4. *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 of Proposition 4* Define $r = (\bar{\pi} - \underline{\pi})/2$. Then $\underline{\pi} = \hat{\pi} - r$ and $\bar{\pi} = \hat{\pi} + r$. Substituting into $p^{**}$ gives

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

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. □

## 4. Model and Characterization

Now that the benefits of the meta-utility theory are clear, we provide a full formal explication of the theory, including a representation result. For this, we employ the following version of Savage's (1954) setting of purely subjective uncertainty.[10] The set of final outcomes is a non-degenerate interval of prizes $X = [\ell, m] \subset \mathbb{R}$. Uncertainty is described by a set $\Omega$ of states of the world and an algebra $\Sigma$ of subsets of $\Omega$ that we refer to as events. For any pair of events $A, B \in \Sigma$, $B \backslash A$ denotes the relative complement of $A$ with respect to $B$.

The objects of choice, denoted by $f \in \mathcal{F}$, are simple acts, that is, mappings $f : \Omega \to X$ that are measurable with respect to $\Sigma$ and have finite range. We identify any $x \in X$ with the constant act that yields $x$ no matter which state $\omega$ in $\Omega$ obtains. For any pair of acts $f$ and $g$ 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.

We begin our description of the family of Ordinal Hurwicz Expected Utility (OHEU) maximizers with the notion of a prior (belief). Intuitively, a prior is a probability measure over a collection $\mathcal{R} \subseteq \Sigma$ of (risky) events. We generalize from the most familiar approach in that these risky events need not form an algebra. Although our approach would hold in the special case in which $\mathcal{R}$ is an ($\sigma$-)algebra, by allowing $\mathcal{R}$ to be a more general collection, we can avoid a concern that Luce (2000, p. 15) and Abdellaoui and Wakker (2005) raise. In short, requiring $\mathcal{R}$ to be an algebra places a potentially extreme demand on the DM, since it requires them to assign probabilities to the joint occurrence of arbitrary events, although the DM may not have any particular expectations regarding these joint events.

The example in Luce (2000) illustrates: suppose a DM is deciding on how to travel from New York to Boston, and is considering the risk of delay if going by air, bus, car, or train. The DM can evaluate the risk of delay for each of these modes of travel separately, without needing to know the joint distributions of the delay risks (e.g., weather might delay flights and disrupt traffic). However, it is not always obvious how these types of disruptions correlate (see also Seidenfeld and Wasserman 1993, for an example where knowledge of marginal distributions is more useful to a DM than information about an ambiguous joint distribution). Luce reckons that there are at least 10,000 states of nature representing possible combinations of levels of delays. It is easy to see that there would be a combinatorial explosion in the number of possibilities the DM would need to consider if having to assign probabilities to all joint distributions. Knowing the marginal distributions is clearly enough. By allowing the collection of risky events to be weaker than an algebra, we can consider the DM as still facing risk, and not ambiguity, even if the DM does not know all the joint distributions.

[10] Our results also go through in an Anscombe and Aumann (1963) framework.

As pointed out in Abdellaoui and Wakker (2005), it is possible to avoid difficulties like the one Luce raises by following the approach of Kopylov (2007), and requiring only that the collection of risky events $\mathcal{R}$ be a *mosaic*.[11] Kopylov (2007) defines a mosaic as a collection satisfying the following three conditions:

M.1 $\Omega \in \mathcal{R}$;

M.2 If $R \in \mathcal{R}$ then $\Omega \backslash R \in \mathcal{R}$;

M.3 If $R_1, \ldots, R_n$ are all in $\mathcal{R}$ and form a partition of $\Omega$, then for all $i, j \in \{1, \ldots, n\}$, $R_i \cup R_j \in \mathcal{R}$.

Conditions M.1 and M.2 are standard and guarantee that trivial and impossible events are $\mathcal{R}$-measurable, and more generally that if it is possible to assign a probability to the occurrence of some event $E$, then it is also possible to assign a probability to $E$ not occurring. Condition M.3 generalizes the usual condition of risky events forming an algebra. It requires only that the coarsening of any risky partition is also risky. We refer to an act $f$ as risky if it is $\mathcal{R}$-measurable with respect to $\mathcal{R}$. All other acts are ambiguous.

As $\mathcal{R}$ need not be closed under intersections, unions, or even disjoint unions, we specify the events for which a prior is sufficiently fine-grained. Roughly speaking, we call a prior *finely ranged* over some event $E \in \mathcal{R}$ if there are risky subsets of $E$ with arbitrarily small probability. Following Kopylov (2007), we extend this idea to collections of events. Doing so is simplified by introducing some notation.

Given a finite collection of risky events $\mathcal{E} \subseteq \mathcal{R}$, set

$$\mathcal{R} \Cap \mathcal{E} := \{R \in \mathcal{R} \colon (\forall E \in \mathcal{E}) R \cap E \in \mathcal{R}\} .$$

That is, an event $R$ is in $\mathcal{R} \Cap \mathcal{E}$ if for every (risky) event $E$ in $\mathcal{E}$ its intersection $R \cap E$ is also risky. We let $\mathcal{G} \Cap \mathcal{E}$ denote the set of acts measurable with respect to $\mathcal{R} \Cap \mathcal{E}$.

We think of a prior $\pi$ as finely ranged with respect to $\mathcal{E}$ if $\mathcal{R} \Cap \mathcal{E}$ includes partitions that split the members of $\mathcal{E}$ into arbitrarily small subsets, that is, to subsets all with probability less than any $\varepsilon > 0$ we choose. If $\pi$ is finely ranged with respect to every finite collection of risky events, we call $\pi$ *finely ranged*. That is, $\pi$ is finely ranged if, for all finite $\mathcal{E} \subseteq \mathcal{R}$ and for every $\varepsilon > 0$, we can find a partition $\{R_1, \ldots, R_n\} \subseteq \mathcal{R} \Cap \mathcal{E}$ such that, for each $i \in \{1, \ldots, n\}$, we have $\pi(R_i) < \varepsilon$.

It is natural to view the DM's preferences over $\mathcal{G}$ as reflecting her risk preferences since the prior $\pi$ can be used to map each (risky) act in $\mathcal{G}$ to a lottery. We do so by defining a function $L_g := \pi \circ g^{-1}$, so that for each prize $x \in X$, we think of $L_g(x)$ as the probability of receiving prize $x$. We can then interchangeably discuss the act $g$ or the associated lottery $L_g$, as convenient. If $g$ is clear from context, we simply write $L$.

[11] He's (2021) two-stage evaluation model of decision under ambiguity also treats risky events as forming a mosaic.

The standard model for risk preferences is (subjective) expected utility, which in addition to a prior $\pi$ employs a Bernoulli utility $v: X \longrightarrow \mathbb{R}$ to represent the DM's preferences over $\mathcal{G}$ by the function $\sum_{x\in X} \pi(g^{-1}(x))v(x) = \sum_{x\in X} L(x)v(x)$.

Our model extends preferences from the risky acts in $\mathcal{G}$ to all simple acts $\mathcal{F}$, including ambiguous ones, using the expected utilities of risky acts as a building block. We accomplish this using a *meta-utility aggregator* (or just a meta-utility), which is a function that maps pairs of expected utility representations onto the range $v(X)$ of possible utilities of prizes in $X$. That is, a meta-utility aggregator is a continuous map

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

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

We think of $U$ as taking a lower bound and an upper bound on the utility of a given act $f \in \mathcal{F}$, and providing a numerical representation, which has the familiar properties of an ordinal utility function.[12] In the special case of a risky act $g \in \mathcal{G}$, we take the expected utility of $g$ as both the lower and upper bound, and it is natural to require that the meta-utility leaves this representation unchanged; i.e. we stipulate that $U(u,u) = u$, hence selecting a unique representation.

We now have all the ingredients required to define an Ordinal Hurwicz Expected Utility.

Definition 2 (OHEU Representation). The function $W\colon \mathcal{F} \to \mathbb{R}$ is an *Ordinal Hurwicz Expected Utility* (OHEU) if there exists a triple $\langle \pi\,, v\,, U\rangle$, where $\pi$ is a finely ranged prior defined on a mosaic of events $\mathcal{R}\,(\subseteq \Sigma)$, $v$ is a Bernoulli utility, and $U$ is a meta-utility aggregator, such that, letting

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

we have

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

We say a function $V\colon \mathcal{F} \to \mathbb{R}$ represents the binary relation $\succsim$ if for all $f\,, f' \in \mathcal{F}$, we have $f \succsim f'$ if and only if $V(f) \geqslant V(f')$. Theorem 1, below, is a representation theorem for an OHEU. The

[12] This feature of our representation reflects an important difference between the meta-utility theory and Gul and Pesendorfer's (2014, 2015) expected uncertain utility theory. Essentially, the meta-utility theory takes expectations (of risky acts) before considering ambiguity, while Gul and Pesendorfer's theory implements these operations in the reverse order, taking the expectation second. One consequence of this is that their utility representation for ambiguous preferences is cardinal, assimilating utility under ambiguity to utility under risk. In the special case of $\alpha$-MEU preferences, the approaches coincide because the aggregation is linear, so the order of taking expectations does not matter, as we remark above informally in Footnote 2.

axiomatization draws on Kopylov (2007) but in order to state the axioms we require a couple more definitions.

An event $A$ is *null* if for all pairs of acts $f$ and $g$, it is the case that $f_A g \sim g$. These will be the events the DM's prior assigns probability zero.

Null events are trivially risky. We follow Zhang (2002) and Kopylov (2007) and say that, in general, events are risky if they satisfy a separability property, which is weaker than the sure-thing principle. Using Zhang's (2002) terminology, we say an event $E$ is *c-separable* if, for all pairs of outcomes $x, y \in X$ and all pairs of acts $f, g \in \mathcal{F}$,

$$x_E f \succsim x_E g \Longrightarrow y_E f \succsim y_E g, \text{ and}$$
$$f_E x \succsim g_E x \Longrightarrow f_E y \succsim g_E y$$

In words, if the DM views an event (and hence, its complement) as risky, then whether they prefer act $f$ or $g$ conditional on that event (or on its non-occurrence) is independent of the prize received in case the event does not (or does) occur. The comparison is required to hold only if the alternative to a given act is a constant prize, in order to allow for the possibility that subsets of subjectively risky acts may be subjectively ambiguous (see Kopylov 2007, p. 245, on this point). We denote the set of c-separable events by $\mathcal{R}_c$, the non-null c-separable events by $\mathcal{R}_c^+$, and the c-separable acts by $\mathcal{G}_c$.

The following six axioms for binary relations on $\mathcal{F}$ are necessary and sufficient for there to exist an OHEU representation. The first three are standard: the usual ordering axiom, followed by a monotonicity requirement that seems very natural in our setting where acts yield monetary prizes, and a continuity axiom. Taken together, these three axioms ensure the certainty equivalent of each act is well-defined.[13]

Ordering The binary relation $\succsim$ is complete and transitive.

Strict Statewise Dominance For all pairs of acts $f$ and $h$ in $\mathcal{F}$, $f \gg h \Longrightarrow f \succ h$.

Outcome Continuity For any sequence of acts $\{f_n\}_{n=1}^{\infty}$ with each $f_n \in \mathcal{F}$ and any pair of acts $g$ and $h$ for which, for all $n$, $g \succsim f_n \succsim h$, if $\{f_n\}$ converges uniformly to $f$ then $g \succsim f \succsim h$.

The next two axioms are adapted from Kopylov (2007) and correspond to Savage's postulates **P4** and **P6**, respectively, restricted to apply only to risky events and acts.

C-separable Comparative Probability For any pair of c-separable events $A$ and $B$ in $\mathcal{R}_c$, and any four outcomes, $x > y$ and $w > z$, $x_A y \succsim x_B y$ implies $w_A z \succsim w_B z$.

This fourth axiom requires that for any pair of risky events, a preference to bet on one event over another should only depend on the subjective relative likelihood of the two events and thus should not depend on the stakes.

[13] For a more general discussion of certainty equivalents and ambiguity, see Grant et al. (2022).

C-separable Event Continuity For any finite collection of c-separable events $\mathcal{E} \subset \mathcal{R}_c$, and for any pair of $\mathcal{E}$-measurable acts $f \succ g$, there exists a finite partition $\{R_1, \ldots, R_k\}$ of $\Omega$ such that, for all $i \in \{1, \ldots, k\}$, $R_i \subset \mathcal{R} \Cap \mathcal{E}$ and $\ell_{R_i} f \succ m_{R_i} g$.

Recall that the worst prize in $X$ is $\ell$ and the best prize is $m$. This axiom says that the DM can divide risky events finely enough that, if the DM strictly prefers act $f$ to act $g$, then they also strictly prefer receiving act $f$ in all but a sufficiently unlikely event to receiving act $g$ in all but the same sufficiently unlikely event, regardless of what they receive if that unlikely event were to occur. Intuitively, the DM would cross the street to receive a treasure chest filled with gold, even if there is a small chance of getting killed from crossing the street.

The final axiom, which is novel and implements the key idea behind the meta-utility theory, formalizes the following intuition. To evaluate an arbitrary act $f$, the DM first considers the set of risky acts that $f$ strictly statewise dominates, evaluates their certainty equivalents, and takes the supremum as a lower bound for the certainty equivalent of $f$. Next, the DM considers the set of risky acts that strictly 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 the DM associates with the act $f$. The next axiom implies that these intervals fully characterize their (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 act $f'$, then the DM cannot strictly prefer $f'$ to $f$.[14]

C-bounds Monotonicity Fix a pair of acts $f, f' \in \mathcal{F}$. If, for each outcome $x \in X$,

$$(\forall g \in \mathcal{G}_c : g \leqslant f)\, x \succsim g \quad \Longrightarrow \quad (\forall g' \in \mathcal{G}_c : g' \leqslant f')\, x \succsim g'; \text{ and}$$
$$(\forall g' \in \mathcal{G}_c : g' \geqslant f')\, g' \succsim x \quad \Longrightarrow \quad (\forall g \in \mathcal{G}_c : g \geqslant f)\, g \succsim x,$$

Then $f \succsim f'$.

Our representation result follows. The proof is in the appendix.

Theorem 1. *A binary relation $\succsim$ satisfies* Ordering, Strict Statewise Dominance, Outcome Continuity, C-Separable Comparative Probability, C-Separable Event Continuity, *and* C-Bounds Monotonicity *iff it admits an OHEU representation $\langle \pi, v, U \rangle$, with $\mathcal{R}_c$ as the domain of $\pi$.*

*Furthermore, any two OHEU functions characterized by the pair of triples $\langle \pi, v, U \rangle$ and $\left\langle \hat{\pi}, \hat{v}, \hat{U} \right\rangle$ represent the same preference relation if and only if there exist $a > 0$ 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*

$$(\hat{u}_1, \hat{u}_2) \in \left\{(w, w') \in v(X) \times v(X) : \exists\ f \in \mathcal{F} \text{ with } w = \underline{v}_\pi^f \text{ and } w' = \bar{v}_\pi^f\right\}.$$

[14] Although their formal statements differ, the security-potential dominance axiom of Frick et al. (2022) can be similarly motivated; see also the discussion in Kopylov (2009).

## 5. Conclusion

There is widespread agreement that ambiguity is present in economically important situations, and that it affects decision making. But there is much less agreement on exactly how agents respond to ambiguity—or so you may have been told.

On a closer inspection, the main theories of decision under ambiguity are much more closely related than they appear. There is a key to understanding them, seeing the relationships between them, and extending them: intermediate microeconomics. If you understand expected utility, and if you understand standard examples of two-good ordinal utilities, then you can readily understand and apply the main theories of decision under ambiguity. Similarly, if you understand marginal rates of substitution, then you can easily understand and work with applications of ambiguity attitude.

## Appendix. Proof of Theorem 1

**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 Ordering, Strict Statewise Dominance, and Outcome Continuity are necessary. The necessity of Axioms C-Separable Comparative Probability and C-Separable Event Continuity follows from the necessity proof of Theorem 4.1 in Kopylov (2007, A.7, 261).

To see that C-Bounds Monotonicity also holds for this preference relation, let $\mathcal{R}$ denote the domain of $\pi$ and $\mathcal{G}$ the set of acts that are measurable with respect to $\mathcal{R}$. Fix a pair of acts $f$ and $f'$. 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

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

Hence,

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

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

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

Hence $f\succsim f'$ which is what Axiom C-Bounds Monotonicity delivers.

**Sufficiency.**

We first show that the weak version of STRICT STATEWISE DOMINANCE holds.

LEMMA 1. $f \geqslant h \implies f \succsim h$.

*Proof* Consider the pair of sequences $f_n = \frac{1}{n}m + \left(\frac{n-1}{n}\right) f$ and $h_n = \frac{1}{n}\ell + \left(\frac{n-1}{n}\right) h$. By construction $f_n \gg h_n$, and hence, STRICT STATEWISE DOMINANCE, $f_n \succ h_n$. Since $f_n$ (respectively, $h_n$) converges uniformly to $f$ (respectively, $h$), OUTCOME CONTINUITY implies $f \succsim h$, as required. □

We next establish that the certainty equivalent of every act is well-defined.

LEMMA 2. *For any $f \in \mathcal{F}$ there is a unique $x \in X$ such that $x \sim f$.*

*Proof* We have $m \succsim f \succsim \ell$ by Lemma 1. Suppose by way of contradiction there exists $x^* \in (\ell, m]$ such that $x^* \succ f$ and $f \succ (x^* - \varepsilon)$ for small positive $\varepsilon$. Since $(x^* - \varepsilon) \to x^*$ as $\varepsilon$ goes to 0, OUTCOME CONTINUITY implies that $f \succsim x^*$, contradicting $x^* \succ f$. Thus there exists an outcome $x \in X$ such that $x \sim f$, and by STRICT STATEWISE DOMINANCE it is unique. □

It follows from Theorem 5.1(I) of Kopylov (2007, 246) that $\mathcal{R}_c$ is a mosaic. Our next step is to show the hypothesis of Theorem 5.2(II) holds since we can then invoke this theorem to establish the restriction of $\succsim$ to $\mathcal{G}_c$ admits a subjective expected utility representation characterized by a prior $\pi$ defined on $\mathcal{G}_c$ and a Bernoulli index $v$.

The properties that Theorem 5.2(II) requires for $\succsim$ to satisfy are Savage's postulates **P1**, **P3**, **P4**, **P5** and **P6** restricted to apply only to c-separable events and acts. From our axioms we see this is immediate for all of them except **P3**. The next lemma, however, shows that this property holds as well.

LEMMA 3. *If $A \in \mathcal{R}_c^+$ and $y > x$ then $y_A h \succ x_A h$ for all $h \in \mathcal{G}_c \Cap \{\Omega \backslash A\}$.*

*Proof* Consider $A \in \mathcal{R}_c^+$, $h \in \mathcal{G}_c \Cap \{\Omega \backslash A\}$ and $x < y$. Then, there exist $f, g \in \mathcal{G}_c \Cap \{A\}$ and $h'$ in $\mathcal{G}_c \Cap \{\Omega \backslash A\}$ such that $f_A h' \succ g_A h'$, which implies $m_A h \succsim f_A h' \succ g_A h' \succsim \ell_A h'$ by Lemma 1. C-separability implies $m \succ \ell_A m$. Applying C-SEPARABLE COMPARATIVE PROBABILITY yields $y \succ x_A y$ which again by c-separability implies $y_A h \succ x_A h$ as desired. □

Let $\langle \pi, v \rangle$ characterize the subjective expected utility representation of $\succsim$ restricted to $\mathcal{G}_c$.

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

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

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

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

defined by setting $\hat{U}(w, w') := v(x)$, where $x$ is the certainty equivalent of an act $f \in \mathcal{F}$ for which $w = \underline{v}_\pi^f$, $w' = \bar{v}_\pi^f$. To show $\hat{U}(\cdot, \cdot)$ 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 $(\underline{v}^f_\pi, \bar{v}^f_\pi) = (w, w')$ and $\left(\underline{v}^{\hat{f}}_\pi, \bar{v}^{\hat{f}}_\pi\right) = (\hat{w}, \hat{w}')$. The two (in)equalities $\underline{v}^f_\pi = (\geqslant)\ \underline{v}^{\hat{f}}_\pi$ and $\bar{v}^f_\pi = (\geqslant)\ \bar{v}^{\hat{f}}_\pi$ imply for any outcome $x \in X$ :

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

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

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

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

So by C-Bounds Monotonicity it follows $f \sim (\succsim)\ \hat{f}$, as required.

Finally, take $U : \{(w, w') \in v(X) \times v(X) : w \leqslant w'\} \longrightarrow \mathbb{R}$, 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 $\succsim$. □

## Acknowledgments

An earlier version of this manuscript was titled "Worst- and Best-Case Expected Utility and Ordinal Meta-Utility." Thanks to Efsthathis Avdis, Gerrit Bauch, Ken Binmore, Robin Cubitt, John Hey, Mamoru Kaneko, M. Ali Khan, Mitri Kitti, László Kóczy, Michał Lewandowski, Alan Martina, Brendan Pass, Paul Pedersen, Marcus Pivato, Frank Riedel, Hendrik Rommeswinkel, Matthew Ryan, Hannu Salonen, Larry Samuelson, Timothy Shields, Emilson Silva, Marie-Louise Vierø, participants at AETW 2019, FUR 2018, SAET 2018, the University of Alberta 2019 Workshop on Pressing and Challenging Changes in Economics, the University of Grenoble Alpes / GAEL 2019 JE on Ambiguity and Strategic Interaction, the 2020 World Congress of the Econometric Society, the 6th World Congress of the Game Theory Society, the workshops at University of Nottingham and University of Turku, and the Bayreuth-MCMP Philosophy and Decision Theory workshop. Patricia Rich gratefully acknowledges support through the German Research Foundation project number 315078566 ("Knowledge and Decision"). Part of this research was undertaken during Jack Stecher's visits to Australian National University in 2018 and 2019 and Patricia Rich's visit to Australian National University in 2019, and they gratefully acknowledge the hospitality and support they received during their stays. Corresponding author: Simon Grant.

### References

Abdellaoui M, Wakker PP (2005) The likelihood method for decision under uncertainty. *Theory and Decision* 58(1):3–76.

Abdellaoui M, Zank H (2023) Source and rank-dependent utility. *Economic Theory* 75(4):949–981.

Ahn D, Choi S, Gale D, Kariv S (2014) Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics* 5(2):195–223.

Anscombe FJ, Aumann RJ (1963) A definition of subjective probability. *The Annals of Mathematics and Statistics* 34:199–205.

Arrow KJ, Hurwicz L (1972) An optimality criterion for decision-making under ignorance. Carter C, Ford J, eds., *Uncertainty and Expectations in Economics*, volume 1 (Oxford, England: Basil Blackwell & Mott Ltd.).

Asparouhova E, Bossaerts P, Eguia J, Zame W (2015) Asset pricing and asymmetric reasoning. *Journal of Political Economy* 123(1):66–122.

Baillon A, Driesen B, Wakker PP (2012) Relative concave utility for risk and ambiguity. *Games and Economic Behavior* 75(2):481–489.

Becker SW, Brownson FO (1964) What price ambiguity? or the role of ambiguity in decision-making. *Journal of Political Economy* 72(1):62–73.

Ben-Porath E, Dekel E, Lipman BL (2014) Optimal allocation with costly verification. *American Economic Review* 104(12):3779–3813.

Berger L, Emmerling J, Tavoni M (2017) Managing catastrophic climate risks under model uncertainty aversion. *Management Science* 63(3):749–765.

Bernanke B, Gertler M, Gilchrist S (1996) The financial accelerator and the flight to quality. *Review of Economic Studies* 78(1):1–15.

Binmore KG (2009) *Rational Decisions* (Princeton, NJ: Princeton University Press).

Bossaerts P, Ghirardato P, Guarnaschelli S, Zame W (2010) Ambiguity in asset markets: Theory and experiments. *Review of Financial Studies* 23(4):1325–1359.

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

De Castro L, Yannelis NC (2018) Uncertainty, efficiency and incentive compatibility: Ambiguity solves the conflict between efficiency and incentive compatibility. *Journal of Economic Theory* 177:678–707.

Einhorn HJ, Hogarth RM (1986) Decision making under ambiguity. *Journal of Business* 59(4, Part 2):S225–S250.

Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75(4):643–669.

Etner J, Jeleva M, Tallon JM (2012) Decision theory under ambiguity. *Journal of Economic Surveys* 26(2):234–270.

Frick M, Iijima R, Le Yaouanq Y (2022) Objective rationality foundations for (dynamic) $\alpha$-meu. *Journal of Economic Theory* 200:105394.

Gale D, Hellwig M (1985) Incentive-compatible debt contracts: The one-period problem. *Review of Economic Studies* 52(4):647–663.

Ghirardato P, Maccheroni F, Marinacci M (2004) Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118(2):133–173.

Ghirardato P, Marinacci M (2002) Ambiguity made precise: A comparative foundation. *Journal of Economic Theory* 102:251–289.

Ghirardato P, Siniscalchi M (2012) Ambiguity in the small and in the large. *Econometrica* 80(6):2827–2847.

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

Good IJ (1966) Symposium on current views of subjective probability: Subjective probability as the measure of a non-measurable set. *Studies in Logic and the Foundations of Mathematics*, volume 44, 319–329 (Elsevier).

Good IJ (1983) *Good Thinking: The Foundations of Probability and Its Applications* (University of Minnesota Press).

Grant S, Polak B (2013) Mean-dispersion preferences and constant absolute uncertainty aversion. *Journal of Economic Theory* 148(4):1361–1398.

Grant S, Rich P, Stecher J (2022) Bayes and Hurwicz without Bernoulli. *Journal of Economic Theory* 199:Article Number 105027.

Gul F, Pesendorfer W (2014) Expected uncertainty utility. *Econometrica* 82(1):1–39.

Gul F, Pesendorfer W (2015) Hurwicz expected utility and subjective sources. *Journal of Economic Theory* 159:465–488.

Hansen LP (2014) Nobel lecture: Uncertainty outside and inside economic models. *Journal of Political Economy* 122(5):945–987.

Hansen LP, Sargent TJ (2001) Robust control and model uncertainty. *American Economic Review* 91(2):60–66.

He Y (2021) Revisiting Ellsberg's and Machina's paraoxes: A two-stage evaluation model under ambiguity. *Management Science* 67(11):6897–6914.

Hurwicz L (1951) Optimality criteria for decision making under ignorance. Discussion Paper, Statistics 370, Cowles Commission.

Izhakian Y (2017) Expected utility with uncertain probabilities theory. *Journal of Mathematical Economics* 69:91–103.

Jaffray JY, Wakker PP (1993) Decision making with belief functions: Compatibility and incompatibility with the sure-thing principle. *Journal of Risk and Uncertainty* 7(3):255–271.

Kadane JB (1992) Healthy scepticism as an expected-utility explanation of the phenomena of Allais and Ellsberg. *Theory and Decision* 32(1):57–64.

Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. *Econometrica* 73(6):1849–1892.

Kopylov I (2007) Subjective probabilities on "small" domains. *Journal of Economic Theory* 133:236–265.

Kopylov I (2009) Choice deferral and ambiguity aversion. *Theoretical Economics* 4(2):199–225.

Krasa S, Villamil AP (2000) Optimal contracts when enforcement is a decision variable. *Econometrica* 68(1):119–134.

Luce RD (2000) *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches* (Mahwah, NJ: Lawrence Erlbaum Associates).

Maccheroni F, Marinacci M, Rustichini A (2006) Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6):1447–1498.

Machina MJ (2014) Ambiguity aversion with three or more outcomes. *American Economic Review* 104(12):3814–3840.

Manski CF (1981) Learning and decision making when subjective probabilities have subjective domains. *The Annals of Statistics* 9(1):59–65.

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

Pintér M (2022) How to make ambiguous strategies. *Journal of Economic Theory* 202:Article Number 105459.

Reinganum JF, Wilde LL (1986) Equilibrium verification and reporting policies in a model of tax compliance. *International Economic Review* 27(3):739–760.

Rigotti L, Shannon C, Strzalecki T (2008) Subjective beliefs and ex ante trade. *Econometrica* 76(5):1167–1190.

Rothschild M, Stiglitz JE (1970) Increasing risk: I. a definition. *Journal of Economic Theory* 2(3):225–243.

Savage LJ (1954) *The Foundations of Statistics* (John Wiley & Sons).

Schlag K, Zapechelnyuk A (2020) Compromise, don't optimize: Generalizing Perfect Bayesian Equilibrium to games with ambiguity. Unpublished manuscript, University of Vienna and University of St. Andrews.

Seidenfeld T, Wasserman L (1993) Dilation for sets of probabilities. *The Annals of Statistics* 21(3):1139–1154.

Siniscalchi M (2009) Vector expected utility and attitudes toward variation. *Econometrica* 77(3):801–855.

Suppes P (1974) The measurement of belief. *Journal of the Royal Statistical Society, Series B (Methodological)* 36(2):160–191.

Townsend RM (1979) Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* 21(2):265–293.

Vierø ML (2012) Contracting in vague environments. *American Economic Journal: Microeconomics* 4(2):104–130.

Wald A (1950) *Statistical Decision Functions* (New York: Wiley).

Wang F (2022) Rank-dependent utility under multiple priors. *Management Science* 68(11):8166–8183.

Zhang J (2002) Subjective ambiguity, expected utility and Choquet expected utility. *Economic Theory* 20(1):159–181.