

Source-dependent Uncertainty Aversion

Madhav Chandrasekher*

Current Draft: May 9, 2017

Abstract: Uncertainty aversion is the name given to the fact that decision-makers (DMs), facing payoff uncertainty, are sometimes willing to trade off mean value for a reduction in uncertainty. This paper studies the phenomenon of *source dependent uncertainty aversion*. This is the name given to the observation that aversion to uncertainty is dependent on the source of uncertainty. The paper introduces, axiomatizes, and identifies a model of source-dependent uncertainty aversion. The model accommodates the [Ellsberg \(1961\)](#) evidence on uncertainty aversion and the [Fox and Tversky \(1995\)](#) evidence on source-dependent betting preferences.

1 Introduction

Ellsberg’s famous thought experiment ([Ellsberg \(1961\)](#)) suggested that decision-maker’s confronted with vague probabilities will frequently select a bet with a lower, but sure, payoff, over bets which offer potentially larger rewards but are also more uncertain. This phenomenon is often referred to as uncertainty aversion. A related concept is the notion of *source-dependence*. Source-dependence refers to the observation that, *ceteris paribus*, a decision-maker’s sensitivity to uncertainty is dependent on the underlying source of uncertainty. For example, consider two asset classes with equally uncertain return but where the source of uncertainty for one asset is different than the other, e.g. a familiar example (and an illustration of the well-known *home-bias puzzle*) is betting on stock markets from two different countries. Source preference, in this case, denotes a preference for one asset class over the other.

This paper presents a new model of source-dependence whose intention is to accommodate both source-dependence and uncertainty aversion. That is, we model decision-makers who not only perceive uncertainty and are averse to its presence, but whose aversion varies with the underlying mechanism which generates the uncertainty (borrowing language from [Abdellaoui et al. \(2011\)](#)) – e.g. in the home bias puzzle, the mechanisms are the country-dependent factors which affect the prices of local stocks. These mechanisms are more commonly referred as “sources of uncertainty” and betting preference are said to exhibit source-dependence when some particular feature of these preferences, e.g. conformity with the expected utility hypothesis, varies specifically with the source. Existing work on source-dependent betting preferences imposes one of two hypotheses: (i) *within*-source betting preferences

*Current e-mail address: mchandrasekher@ucdavis.edu.

satisfy expected utility, and (ii) sources of uncertainty are objective, meaning that the bets facing the DM are observably framed as bets on different event classes (e.g. [Ergin and Gul \(2009\)](#)).¹ In this paper, we relax these hypothesis. First, within-source preferences are not required to conform to expected utility. Second, the sources of uncertainty are endogenously derived from the DM’s betting preferences, i.e. bets are not observably framed as acts on different event classes. In so doing, we obtain a relatively general model of uncertainty aversion which can accommodate a wide variety of uncertainty (ambiguity) averse behavior. Importantly, our model explains source-dependent betting preferences as described in [Fox and Tversky \(1995\)](#) and [Heath and Tversky \(1991\)](#) (see example 1, section 2), and reconciles this behavior within a model of uncertainty aversion. Source-dependence without uncertainty suggests that decision-makers reject some bets over others purely on grounds of familiarity with the events on which bets are placed, even when given better (and known) odds on less familiar bets. Arguably, a familiarity-based preference for bets should extend to situations where odds are uncertain. In this case, analogy suggests that a decision-maker (DM) averse to uncertainty over equally familiar and uncertain bets might nevertheless violate uncertainty aversion when comparing two bets which are equally uncertain, but not equally familiar.

While uncertainty aversion and source-dependence have been well-examined in separation, the concept of source-dependent uncertainty aversion has received less attention.² Perhaps one reason for this is that, as noted in [Fox and Tversky \(1995\)](#), source-dependence predicts behavior which can be at odds with uncertainty (ambiguity) aversion. The quote below is their summary of [Heath and Tversky \(1991\)](#)’s *competence hypothesis*, which is a general name [Heath and Tversky \(1991\)](#) give to the finding that, *ceteris paribus*, people seem willing to take on bets on familiar objects (e.g. sports) over those on unfamiliar ones (e.g. politics).

Ambiguity aversion implies a preference for chance over uncertainty because the probabilities associated . . . are necessarily vague or imprecise. In contrast, the competence hypothesis predicts, for example, that a sports fan will prefer to bet on the game than on chance.

A chance event here is the name given to a bet with objectively known odds, e.g. the archetypal chance event being the outcome of a toss of an unbiased coin.

¹The latter restriction is not common to all models of source-dependence. For example, while it is used in [Ergin and Gul \(2009\)](#) (e.g. choice objects are acts on a product state space), the event classes which constitute separate sources are elicited in [Abdellaoui et al. \(2011\)](#). Also, [Gul and Pesendorfer \(2015\)](#) define a source to be a pair (\mathcal{F}, ν) where ν is a measure and \mathcal{F} is the (maximal) domain on which the measure lives. In this model, different sources correspond to different extensions of some base measure ν_0 with domain \mathcal{F}_0 .

²As we discuss in detail later, [Gul and Pesendorfer \(2015\)](#)’s model of source-dependent subjective expected utility is the closest.

The [Heath and Tversky \(1991\)](#) findings suggest a difficulty in reconciling source-dependent behavior with ambiguity aversion. Any such reconciliation would seem to necessitate a violation of uncertainty aversion, which is one of the most well-known and ubiquitous axioms in the ambiguity aversion literature.

In order to link together the notion of source-dependence and uncertainty aversion, we will also need to reconsider what it means to be a “source of uncertainty”. In [Fox and Tversky \(1995\)](#) a source of uncertainty is modeled as a class of events. In this paper we think of a source as a set of priors. *Why?* When the collection of events constituting a Fox-Tversky source satisfies the axioms of subjective probability, there is no distinction between thinking of a source as (i) a class of events or (ii) as a quantitative representation of a likelihood relation on this class. Thus, if within-source preferences were to be expected utility we could, equivalently, think of a collection of sources as a collection of probability measures, with one class of events associated to each measure. However, we show by example that we can nest the Fox-Tversky notion of a source-preference within our model, but that doing so requires *sets* of multiple priors as opposed to a collection of priors with one for each event class (as the Fox-Tversky definition would suggest).³ This then takes us to a model which looks like a “multiple sources” version of the classical [Gilboa and Schmeidler \(1989\)](#) multiple priors (MEU) model. Intuitively, our DM’s choose between acts by first determining a set of priors which are relevant to evaluating that act. Two acts are considered “familiar”, in the sense of source-preference, if they have the same set of relevant priors. Uncertainty aversion applies in comparing acts with the same set of relevant priors, but can be violated when comparing acts with distinct sets of relevant priors.

Our model also provides a new perspective on the question of how to separate ambiguity attitude from the presence of ambiguity. To recapitulate this latter issue, recall that the classical MEU model and most of its generalizations are multiple prior models. Instead of a single prior used to evaluate the expected value of a bet, there is a set of priors. And for each act we, therefore, get an associated set of expected values. Most existing models can be distinguished from one another by the manner in which they make a selection from this set. The MEU selection corresponds to the “min” selector – i.e. where the minimum value is always selected from any set. Various alternatives have been studied as well, e.g. two leading examples are the α -MEU model ([Ghirardato et al. \(2004\)](#)) which takes a weighted average between the “max” selector and the “min” selector and the smooth ambiguity model ([Klibanoff et al. \(2005\)](#)) which induces a random selection from the set of expected values.

³[Fox and Tversky \(1995\)](#) is consistent with this fact since, both in their paper and [Abdellaoui et al. \(2011\)](#), preferences over bets measurable with respect to distinct event classes are numerically represented by a capacity, as opposed to a single coherent probability.

Under these models, presence of ambiguity is modeled with multiple priors and sensitivity to ambiguity is modeled via the selection map (which maps an act to a selected expected value). Critically, the domain of the selection map is the set of expected values associated to a given act and not the set of priors which induce these expected values. Our model is also a multiple priors model, but in order to model source-dependence the domain of our selection map is the set of priors itself. The reason is that if two acts induce the same set of expected values, then a selector which only “sees” the expected values and pays no attention to the underlying priors cannot distinguish between these acts. Yet, a source-dependent MEU decision-maker may exhibit a preference for one set of expected values over another since he is less averse to uncertainty over those events on which that act yields an uncertain payoff, even when the numerical expected values are the same. In this way, our model obtains a partial separation of presence of ambiguity from sensitivity to ambiguity. It is partial since, conditional on a source (set of relevant priors), we force the DM to be completely averse to uncertainty. However, it also turns out to be the case that there are no observable implications (for choices between acts) from taking an a priori more permissive selector, e.g. the α -maxmin selector of [Ghirardato et al. \(2004\)](#) rather than the MEU selector (see proposition 1).⁴

The remaining sections are organized as follows. Section 2 describes the primitives, axioms, and model(s) and demonstrates that our model accommodates source-preferences. Section 3 presents the representation theorem and the identification result. These are our two main results. The proofs of these results are lengthy. We give a sketch in the text and leave more difficult details to the appendix. Section 4 concludes.

2 Axioms and Model

2.1 Axioms

The primitive in this paper is a binary relation on Anscombe-Aumann acts. Formally, let X denote a finite set and put

$$\Delta(X) := \{p \in \mathbf{R}^{|X|} : \text{(i) } p_i \geq 0, \text{ (ii) } \sum_i p_i = 1\}$$

Let $\mathcal{S} := \{s_1, \dots, s_k\}$ denote a finite set of objective states and put $B(\mathcal{S}) := \Delta(X)^{\mathcal{S}}$, i.e. functions mapping from states to lotteries. The primitive is a preference relation

⁴More precisely, if we replace the “min” selector with a rationalisable choice function then the resulting model can be nested in a generalized version of our model, which does use the min-selector within sources, where there are possibly infinitely many sources.

$\succeq \subseteq B(\mathcal{S}) \times B(\mathcal{S})$. We impose the following axioms on \succeq .

Axiom 1: (Order) \succeq is complete and transitive.

Axiom 2: (Continuity) For each $f \in B(\mathcal{S})$ the sets $\{g \in B(\mathcal{S}) : f \succeq g\}, \{g \in B(\mathcal{S}) : f \preceq g\}$ are closed.

Axiom 3: (c-Independence) For $\ell \in \Delta(X)$ let $\vec{\ell}$ denote the act which gives lottery ℓ in every state. Then, $f \succeq g$ if and only if $\alpha \cdot f + (1 - \alpha) \cdot \vec{\ell} \succeq \alpha \cdot g + (1 - \alpha) \cdot \vec{\ell}, \forall \alpha \in [0, 1]$.

This has been called “strong” c-Independence in [Maccheroni et al. \(2006\)](#), where “weak” c-Independence denotes the direction: $f \succeq g \Rightarrow f\alpha c \succeq g\alpha c$.

Axiom 4: (Monotonicity) If $f(s) \succeq g(s), \forall s \in \mathcal{S}$, then $f \succeq g$.⁵

Axiom 5: (Weak Uncertainty Aversion) For every f there is some ϵ_f such that whenever $g \in B_{\epsilon_f}(f)$ and $g \succeq f$, then $\alpha \cdot g + (1 - \alpha) \cdot f \succeq f, \forall \alpha \in [0, 1]$.

In words, given two topologically close acts between which the DM is indifferent, the DM exhibits a preference for hedging. Compare with the [Gilboa and Schmeidler \(1989\)](#) uncertainty aversion axiom (UA):

UA: For every f, g with $g \succeq f$ we have $\alpha \cdot g + (1 - \alpha) \cdot f \succeq f, \forall \alpha \in [0, 1]$.

Let us recall the motivation for this axiom. Put $g\alpha f := \alpha g + (1 - \alpha)f$. When $g \sim f$, the conclusion that $g\alpha f \succeq f, g$ means that the DM prefers the “hedge” between f, g to either f or g . This is because the mixture of the two acts smooths the payoffs across states where either f is better than g or vice-versa – making the mixed act less risky, albeit with a possibly lower mean. Local uncertainty aversion relaxes this by requiring that a preference for smoothing need only exist among acts which offer similar payoffs in all states, i.e. acts which are topologically close.

2.2 Model

The model in this paper has two pieces: (i) a vNM utility on lotteries, $u : \Delta(X) \rightarrow \mathbf{R}$, and (ii) a finite collection, $\{\Pi_i\}_{i=1}^N$, of closed, convex sets of priors. Using these we define what we call a “local MEU” (after [Gilboa and Schmeidler \(1989\)](#)) functional:

$$(**) U(f) = \max_{i=1, \dots, N} \min_{\pi \in \Pi_i} E_{\pi} u(f).$$

Here $E_{\pi}(\cdot)$ denotes expectation w.r.t. π . Under this model, the value of an act is the maximum of its MEU values taken across the sets Π_i . In this paper, we

⁵Here we abuse notation and let $f(s)$ (resp. $g(s)$) denote the act which gives the constant lottery $f(s)$ ($g(s)$) in every state.

will refer to each set of priors as a source of uncertainty. We claim the local MEU model (***) nests the [Fox and Tversky \(1995\)](#) (resp. [Abdellaoui et al. \(2011\)](#)) notion of a source-preference. Recall their definition of source-dependence (referred to as *source-preference*).

Definition 1. Fix a state space \mathcal{S} and a set of collections of events in this space, \mathcal{E}_i . A betting preference over events in $\cup_i \mathcal{E}_i$ is said to exhibit *source-preference* if it exhibits the following property:

For any fixed $(A, B) \in \mathcal{E}_i \times \mathcal{E}_j$, a DM prefers a bet on A over B and a bet on A^c over B^c .

That is, $\mathbf{1}_A$ (the indicator bet on event A) is preferred to $\mathbf{1}_B$ and $\mathbf{1}_{A^c}$ over $\mathbf{1}_{B^c}$. The local MEU model exhibits (some) ambiguity aversion, but can accommodate this kind of source preference.

Example 1. Fix a state space \mathcal{S} and events $A, B \subseteq \mathcal{S}$, and for simplicity's sake take $A \neq B^c$. Define four measures $\{\pi_i\}_{i=1}^4$ on this space with the following properties:

- i. $\pi_1(A) > \pi_3(A), \pi_2(A) > \pi_1(A), \pi_4(A) > \pi_3(A)$.
- ii. $\pi_2(A^c) > \pi_4(A^c)$.
- iii. $\pi_1(A) > \pi_3(B) > \pi_1(B), \pi_4(B) > \pi_3(B), \pi_2(B) > \pi_1(B)$.
- iv. $\pi_2(A^c) > \pi_4(B^c) > \pi_2(B^c)$.

It is straightforward to verify that for any (finite) state space (with $|\mathcal{S}| \geq 3$) and distinct events $A, B, B \neq A^c$ defined on this space, we can find such a quadruple (π_1, \dots, π_4) . Now take any u and pick two lotteries, ℓ_x, ℓ_y with $u(\ell_x) = x \gg y = u(\ell_y)$. Let $\Pi_1 := \{\pi_1, \pi_2\}, \Pi_2 = \{\pi_3, \pi_4\}$. The local MEU model defined by the pair $(u, \{\Pi_i\}_{i=1}^2)$ gives the source-preference described above. To see this, note that (i), (ii) ensure that the (MEU) value of the act xAy (resp. $xA^c y$) is obtained on the set Π_1 and (iii), (iv) ensure that the value of the act $xB y$ (resp. $xB^c y$) is obtained on the set Π_2 and that the utility of xAy (resp. $xA^c y$) is higher than $xB y$ (resp. $xB^c y$).

As the example suggests, the reason the DM exhibits a source preference (in the Fox-Tversky sense) for event A over event B is that the priors relevant for evaluating event A are more “optimistic” than those that are relevant for evaluating event B . Nevertheless, for two events with the same relevant set of priors the DM evaluates both pessimistically, i.e. in the MEU sense. For this reason, we switch the designation of a “source” from a set of events to the set of priors – implicitly, those priors which are relevant for assessing the value of the binary acts, i.e. xAy , defined on these events. Interestingly, these two ways of thinking about a source of

uncertainty (i.e. as event classes vs. sets of measures on these event classes) turn out to be “essentially” the same. This comes out of our proof of identification for the local MEU model (see theorem 2).⁶

Example 1 cannot, in general, be represented by an MEU model. To see this, assume that we had an MEU representation (u, Π) , so that:

- i. $\min_{\pi \in \Pi} \pi(A) > \min_{\pi \in \Pi} \pi(B)$.
- ii. $\min_{\pi \in \Pi} \pi(A^c) > \min_{\pi \in \Pi} \pi(B^c)$.

Note that (ii) is equivalent to $\min_{\pi \in \Pi} (1 - \pi(A)) > \min_{\pi \in \Pi} (1 - \pi(B))$, which in turn is equivalent to: $\max_{\pi \in \Pi} \pi(A) < \max_{\pi \in \Pi} \pi(B)$. Repeat the argument with (i) to find that it is equivalent to: $\max_{\pi \in \Pi} \pi(A^c) < \max_{\pi \in \Pi} \pi(B^c)$. Hence, an MEU representation of the source-dependent preference is only possible if we have:

- i. $[\min_{\pi} \pi(A), \max_{\pi} \pi(A)] \subseteq [\min_{\pi} \pi(B), \max_{\pi} \pi(B)]$,
- ii. $[\min_{\pi} \pi(A^c), \max_{\pi} \pi(A^c)] \subseteq [\min_{\pi} \pi(B^c), \max_{\pi} \pi(B^c)]$.

Experiments which document source preferences also have results on the bid-ask spreads associated to the bets $\mathbf{1}_A, \mathbf{1}_B$, where A, B come from different sources. For instance, [Fox and Tversky \(1995\)](#) report elicited odds as well as bid-ask prices which give an interval of odds when the bets are subject to uncertainty.⁷ When there is no uncertainty, i.e. known odds, the above conditions require that bet $\mathbf{1}_A$ is exactly indifferent to $\mathbf{1}_B$ (resp. $\mathbf{1}_{A^c}$ is indifferent to $\mathbf{1}_{B^c}$). With uncertainty, the above conditions allow the source preference only if the bet $\mathbf{1}_B$ (resp. $\mathbf{1}_{B^c}$) is more uncertain than the bet $\mathbf{1}_A$ (resp. $\mathbf{1}_{A^c}$) – in the sense that its elicited bid-ask spread is nested within the spread for $\mathbf{1}_B$.⁸ Hence, were we to elicit both source-preferences and bid-ask spreads on bets, then the MEU model cannot explain a preference for $\mathbf{1}_A$ over $\mathbf{1}_B$, where the latter is nevertheless more certain than the former. Now we check that even taking just a source-preference as a primitive (i.e. no elicitation of bid-ask spreads), the Fox-Tversky examples cannot in general be represented by an MEU model.

⁶Fixing a local MEU model, we show identification by showing that the set of “critical” acts, i.e. acts whose strict maximum occurs on some source, uniquely pins down the collections of measures which comprise the local MEU representation. Criticality is *exactly* the property of the bets xAy (resp. xBy) in example 2 above that allowed the construction to go through.

⁷More precisely, [Fox and Tversky \(1995\)](#) infer certainty equivalents and odds (or intervals of odds) are derived from this. In their experimental design, subjects are sequentially offered choices between a bet and a cash alternative. As they note, certainty equivalents are inferred from bid-ask prices: *the certainty equivalent of each prospect was determined by a linear interpolation between the lowest value accepted and the highest value rejected...*, Ibid, pg.273.

⁸As shown in [Ghirardato et al. \(2004\)](#), for a broad class of ambiguity models, the elicited bid-ask spreads pin down, and are in turn determined by, the priors describing the ambiguity models.

Example 2. Let $\mathcal{S} = \{s_1, s_2, s_3\}$, $A = \{s_1\}$, $B = \{s_3\}$, $C = \{s_2\}$ and define three Fox-Tversky sources, $\mathcal{E}_1 = \{A, A^c, \mathcal{S}, \emptyset\}$, $\mathcal{E}_2 = \{B, B^c, \mathcal{S}, \emptyset\}$, $\mathcal{E}_3 = \{C, C^c, \mathcal{S}, \emptyset\}$. Consider the source-preference:

$$\mathbf{1}_A \text{ (resp. } \mathbf{1}_{A^c}) \succ \mathbf{1}_B(\mathbf{1}_{B^c}) \succ \mathbf{1}_C(\mathbf{1}_{C^c}).$$

There is no MEU representation of this preference, though example 1 shows that it does admit a local MEU representation. To see the obstruction, consider a putative MEU representation, say (u, Π) and let $\underline{x} = \min_{\pi \in \Pi} E_{\pi} \mathbf{1}_A$, $\bar{x} = \max_{\pi \in \Pi} E_{\pi} \mathbf{1}_A$, $\underline{z} = \min_{\pi \in \Pi} E_{\pi} \mathbf{1}_C$, $\bar{z} = \max_{\pi \in \Pi} E_{\pi} \mathbf{1}_C$, and similarly define \underline{y}, \bar{y} (for the MEU and max-EU values of $\mathbf{1}_B$). The parameters $(\underline{x}, \bar{x}, \underline{y}, \bar{y}, \underline{z}, \bar{z})$ must satisfy the following inequalities:

1. $\underline{z} < \underline{x}$.
2. $\bar{z} > \bar{x}$.
3. $1 - \bar{x} > \underline{z}$.
4. $1 - \underline{x} > \bar{z}$.
5. $\bar{x} > \underline{x}$.
6. $\bar{x} + \bar{z} > 2\underline{z}$.

A symmetric set of inequalities hold for the quadruples $(\underline{x}, \bar{x}, \underline{y}, \bar{y})$ and $(\underline{y}, \bar{y}, \underline{z}, \bar{z})$.⁹ The first five inequalities follow directly from the source preference. The next uses the source preference and uncertainty aversion: since $\mathbf{1}_A \succ \mathbf{1}_C$ we must have (take a 50:50 mixture, multiply by 2, and note that the MEU functional extends the representation to all scaled multiples of AA acts) $\mathbf{1}_A + \mathbf{1}_C \succeq 2\mathbf{1}_C$. $\bar{x} + \bar{z}$ is an upper bound on the LHS and $2\underline{z}$ is the MEU value of the RHS. We can embed a solution $(\underline{x}, \bar{x}, \underline{y}, \bar{y}, \underline{z}, \bar{z})$ into the solution space of an appropriate LP program. Intuitively, since there six variables and many more (≥ 20) linear restrictions on these variables, the solutions space should be (at least, generically) empty. The formal argument (in

⁹We make two comments. First, there will also be three additional constraints on the sextuplet $(\underline{x}, \bar{x}, \underline{y}, \bar{y}, \underline{z}, \bar{z})$ which do not appear on this list (see the appendix). Second, for a non-MEU example we need three sources. To see this, consider just two Fox-Tversky sources, $\mathcal{E}_1 = \{A, A^c, \mathcal{S}, \emptyset\}$, $\mathcal{E}_2 = \{B, B^c, \mathcal{S}, \emptyset\}$, where $A = \{s_1\}$, $B = \{s_3\}$. There is an MEU representation in this case: let $\pi_1 = (1/6, 1/6, 2/3)$, $\pi_2 = (5/24, 2/3, 1/8)$ and $\Pi = \{\pi_1, \pi_2\}$. The model (u, Π) (for any u with $u(1) > u(0)$) represents the source preference. Note, however, that putting $C = \{s_2\}$ and taking a third source $\mathcal{E}_3 = \{C, C^c, \mathcal{S}, \emptyset\}$ we no longer have a representation. There is, however, a two source local MEU representation. Take $\Pi_1 = \{\pi_1^{\varepsilon}, \pi_2\}$, where π_2 is as above, but $\pi_1^{\varepsilon} = (1/6 + \varepsilon, 1/6 - \varepsilon, 2/3)$. Put $\Pi_2 = \{\pi_3, \pi_4\}$, where $\pi_3 = (1/8 + \varepsilon_1, 2/3, 5/24 - \varepsilon_1)$, $\pi_4 = (5/24 + \varepsilon_2, 1/8, 2/3 - \varepsilon_2)$, and $\varepsilon > \varepsilon_1 > \varepsilon_2$, with $2/3 - \varepsilon_2 > 5/24 - \varepsilon_1$. Then, we have $\mathbf{1}_A, \mathbf{1}_{A^c}$ attain their (local) MEU values on the source Π_1 , whereas $\mathbf{1}_B$ (resp. $\mathbf{1}_{B^c}$) attain their values on the source Π_2 . Moreover, the source preference relations hold, i.e. betting on \mathcal{E}_1 is preferred to betting on either of $\mathcal{E}_2, \mathcal{E}_3$ and betting on \mathcal{E}_2 is preferred to betting on \mathcal{E}_3 .

the appendix) shows that there is an LP program whose solution space contains the solution space to the above inequalities and we check that the alternative to this LP is always non-empty.

Two precedents in the literature which also define sources of uncertainty in terms of subjective beliefs are [Chew and Sagi \(2008\)](#) (CS) and [Gul and Pesendorfer \(2015\)](#) (GP). CS define a source of uncertainty (referred to in their paper as a *conditional small world domain*) as a collection of events subject to some restrictions and show that these restrictions (distinct from the Savage axioms, and based on CS’s concept of exchangeability) are satisfied iff there is a well-defined probability measure representing the likelihood relation on this class of events. The GP model is a little closer to what we do here since it is also a multiple priors model. They define a source as a pair (π, \mathcal{F}_π) consisting of a prior along with a domain of measurable events which is (implicitly) the domain of this prior. This approach ties the presence of uncertainty to a measure extension problem.

To see this, note that under the GP definition of a source these two components – i.e. measures and the domains of definition of the measure – cannot be decoupled. In their model, any two priors with the same domain of definition must agree. When priors are uniquely defined by their (maximal) domain of definition the only way there can be uncertainty (i.e. multiple priors) is if starting from a given prior there is more than one way to extend its domain of definition to a larger class of events. The GP functional form is a special case of the Hurwitz α -maxmin utility and (for finite state spaces) the Hurwitz model turns out to be nested in a version of the local MEU model where we allow for infinitely many sources – see the next sub-section for a discussion. The GP utility is well-defined on any state space, but to derive their results their approach requires a rich, e.g. not finite, state space.

2.3 Comparison to other ambiguity models

Now consider a functional form that allows for a more general selection map than the “min” selector. We let \mathcal{K} denote the set of all closed subsets of \mathbf{R} and require selections from elements of \mathcal{K} to satisfy the following three conditions:

1. $\theta : \mathcal{K} \rightarrow \mathbf{R}$ is continuous w.r.t to the Hausdorff metric on the domain \mathcal{K} .
2. $\theta(A) \in A, \forall A \in \mathcal{K}$.
3. $\theta(B) = \theta(A)$ whenever, $A \subseteq B$ and $\theta(B) \in A$.

The first and third property say that θ comes from maximization of a valuation, say v_θ , on \mathbf{R} . Define a “local θ – MEU” representation via the formula

$$U(f) := \max_{\Pi_i} \theta - \text{MEU}_{\Pi_i}(f)$$

Well-known models of multiple priors use a selection rule which satisfies exactly these or related, e.g. as in (ii), properties.

- i. (**MEU**) Standard selectors such as “min” and “max” satisfy both these conditions; hence, the condition nests both subjective expected utility theory (which is recovered with either the min or max elector) and max-min (MEU) expected utility.
- ii. (**Hurwitz max-min preferences**) We can also relax the hypothesis that θ is necessarily derived from a valuation, v_θ , and nest the well-known model from [Ghirardato et al. \(2004\)](#), i.e. here the utility of an act $U(f)$ is defined using Hurwitz’s α -maxmin criterion:

$$U(f) = \alpha \cdot \min_{\pi \in \Pi} E_\pi u(f) + (1 - \alpha) \cdot \max_{\pi \in \Pi} E_\pi u(f).$$

For example, consider the examples where – for A convex – we put

$$\theta(A) = \alpha \cdot \min(A) + (1 - \alpha) \cdot \max(A).$$

Fixing a pair (u, Π) , consisting of a vNM function on lotteries and a set of priors Π , and applying the selector θ gives the α -MEU model axiomatized in [Ghirardato et al. \(2004\)](#).

One of the most general models of ambiguity aversion is the **variational preferences** representation of [Maccheroni et al. \(2006\)](#). What is the connection between variational preferences and the local MEU model (or, for that matter, the local θ -MEU model)? Recall the primitives of the [Maccheroni et al. \(2006\)](#) model are a pair (u, c) , where u is an affine function on lotteries and c is a function which assigns a cost to each potential prior $\pi \in \Delta(\mathcal{S})$. These are put together to define the following utility on acts:

$$(*) U(f) = \min_{\pi \in \Delta(\mathcal{S})} [E_\pi u(f) + c(\pi)].$$

Since the local θ -MEU model satisfies (strong) c -independence and variational preferences need only satisfy weak c -independence, there is not a nesting of variational preferences within the class of local θ -MEU models.

There is nevertheless an interesting connection. To see this, for each act c define $A_f := \{E_\pi u(f) + c(\pi) : \pi \in \Delta(\mathcal{S})\}$. Then, taking θ to be the choice function induced by the “min” function we recover the variational preference representation of [Maccheroni et al. \(2006\)](#) with this θ on the c -shifted sets A_f . This suggests that changes in ambiguity attitude, as captured by the [Maccheroni et al. \(2006\)](#) model, are different from those suggested by the model in this paper. While there isn’t (as far as we can tell) an clear contrast in functional forms, there is a clear comparison in the (resp.) behavioral characterizations. [Maccheroni et al. \(2006\)](#) show that

the model (*) is characterized by solely the standard axioms (axiom 1, 2, 4), weak c -independence, and the uncertainty aversion axiom (i.e. UA above). By contrast, we will show that the local MEU model is characterized by axioms 1,2,4, strong c -independence, and local uncertainty aversion. Put another way, variational preferences nest the [Gilboa and Schmeidler \(1989\)](#) multiple priors model by weakening c -independence but maintaining uncertainty aversion whereas local MEU preferences maintain strong c -independence but weaken uncertainty aversion.

Now consider a version of the local θ -MEU model where we take the θ -selection map to be the “min” and the collection of priors $\{\Pi_i\}$ to be potentially infinite. Since the collection is potential infinite, we replace the outer “max” with a supremum in the following formula:

$$(**) U(f) = \sup_{\Pi_i} \min_{\pi \in \Pi_i} E_{\pi} u(f).$$

The following proposition shows that the local θ -MEU model is nested within the class of local MEU models, so long as we allow for (potentially) infinitely many sources.¹⁰

Proposition 1. *Every utility U on acts which admits a local θ -MEU representation also has a local MEU representation (with a potentially infinite set of sources).*

The selector map θ is usually interpreted as capturing the DM’s “attitude” towards the presence of uncertainty. For example, in [Ghirardato et al. \(2004\)](#) the α parameter defining the Hurwitz criterion parameterizes the level of uncertainty aversion. Up to the finiteness condition on the “sources”, i.e. the sets of priors, the proposition argues that there is no loss in passing to the “min” selector. As we will see in the proof of theorem 1, the main restriction imposed by local uncertainty aversion (axiom 5) is that it forces the collection of sources to be finite.

2.4 A possible alternative definition of sources

Are there alternative ways to model both sources of uncertainty and uncertainty aversion? One possibility which seems promising and which allows us to go back to modeling sources as collections of events is to add some structure to the [Schmeidler \(1989\)](#) model of Choquet expected utility (CEU). Recall that the source functions $W(\cdot)$ in [Fox and Tversky \(1995\)](#) and [Abdellaoui et al. \(2011\)](#) are taken to be capacities and the CEU model allows for violations of uncertainty aversion – which, as suggested in [Heath and Tversky \(1991\)](#), is required (to some degree) to capture source-preference. We argue now that the CEU approach runs into some difficulties.

¹⁰Essentially the same proof of the proposition shows that the model with the GMM selector can be nested in (**) as well.

Let (u, ν) denote a pair consisting of a vNM utility on lotteries and a capacity $\nu : 2^{\mathcal{S}} \rightarrow [0, 1]$.¹¹

Definition 2. Given a state space \mathcal{S} call any (sigma) algebra of events on \mathcal{S} , \mathcal{E} , a *source*.

Fixing a source \mathcal{E} we let $B_{\mathcal{E}} \subseteq B(\mathcal{S})$ denote the subset of \mathcal{E} -measurable acts. Abusing notation, let $\succeq|_{\mathcal{E}}$ denote the preference restricted to $B_{\mathcal{E}}$ -measurable acts. The umbrella model to which we add structure is the Choquet EU (CEU) model (Schmeidler (1989)) built from the pair (u, ν) .

Definition 3. Fix a collection $\{\mathcal{E}_i\}_{i=1}^N$ of sources. A CEU model (u, ν) is *locally convex* with respect to the collection $\{\mathcal{E}_i\}_{i=1}^N$ if for each i there is a closed, convex set of priors Π_i such that $\nu(E) = \min_{\pi \in \Pi_i} \pi(E)$.

Call this a “local CEU” utility. Let $U(\cdot)$ be the utility on acts generated by the pair (u, ν) . Note that when f is an \mathcal{E}_i -measurable act, then we have the formula $U(f) = \min_{\pi \in \Pi_i} E_{\pi} u(f)$, so that this model bears similarity to the local MEU model. However, the MEU structure holds for only a limited class of acts (for the local CEU model). In particular, any AA act whose associated state-dependent values are all distinct (an open, dense subset of the space of all acts) can only be measurable w.r.t. the power set of the state space \mathcal{S} .¹² Hence, for $U(f)$ (for such an act) to have a local MEU representation it must be the case that one of the algebras \mathcal{E}_i is the entire power set of the state space. Since all acts are measurable w.r.t this algebra, this forces $U(f) = \min_{\pi \in \Pi_i} E_{\pi} u(f)$ for all acts f , i.e. the representation collapses to MEU if we want a generic local MEU representation. If we omit the generic set of acts with distinct state-dependent values, then the local CEU model is not necessarily an MEU but on acts with distinct state-dependent values we cannot express the Choquet integral of f against ν in an MEU form.

The preceding suggests that the CEU structure doesn’t admit a representation where the capacity has a locally convex structure. On the other hand, the CEU model does allow for violations of uncertainty aversion and Schmeidler (1989) shows that this is indicated by non-convexity of the core of the capacity ν . The local MEU model certainly exhibits such non-convexities since the set, $\cup_i \Pi_i$, consisting of the union of priors across all sources need not be convex. This might suggest a nesting of the local MEU model within the CEU class (even in the absence of a locally convex structure of the capacity¹³), but this is not the case. The following example

¹¹Recall that a capacity is just a monotone set function which assigns a non-negative number (wlog between $[0, 1]$) to each subset of states, i.e. event.

¹²Recall that the state space \mathcal{S} is finite.

¹³Note that the two claims, (i) the CEU capacity is locally convex and (ii) the local MEU model has a (u, ν) representation, are not the same.

demonstrates that local MEU models can violate [Schmeidler \(1989\)](#)'s co-monotonic independence axiom.

Example 3. Let $\mathcal{S} = \{s_1, s_2, s_3\}$ and consider two acts f, g defined as follows: (acts are valued in utils)

	s_1	s_2	s_3
$f(s)$	1	$\frac{1}{2} - \varepsilon$	0

	s_1	s_2	s_3
$g(s)$	$\frac{1}{3}$	$\frac{1}{3} - \varepsilon$	0

Where we take $\varepsilon > 0$ to be some very small, positive number. Note that f, g are co-monotonic acts. However, consider the two measures $\pi_1 = (1/3, 1/3, 1/3), \pi_2 = (1/2, 0, 1/2)$, i.e. π_1 is uniform and π_2 puts weight 1/2 on states s_1, s_3 , and note that

$$E_{\pi_1} f < E_{\pi_2} f, E_{\pi_1} g > E_{\pi_2} g,$$

for all small ε . Replace g with $c \cdot g$ and f with $d \cdot f$, where we choose $c, d > 0$ so that

$$(*) E_{\pi_1} c \cdot g = E_{\pi_2} d \cdot f - \varepsilon > E_{\pi_1} d \cdot f,$$

again for some small, positive ε . Consider the local MEU model with sources $\Pi_1 := \{\pi_1\}, \Pi_2 := \{\pi_2\}$. Now take any h with $E_{\pi_1} h > E_{\pi_2} h$ which is comonotone with f, g (i.e. (f, g, h) are pairwise comonotone), so that for all α close to 0 we have

$$E_{\pi_1}(cg)\alpha h > E_{\pi_2}(cg)\alpha h, E_{\pi_1}(df)\alpha h > E_{\pi_2}(df)\alpha h,$$

where $f\alpha g = \alpha f + (1 - \alpha)g$, for generic f, g . By (*), we have

$$E_{\pi_1}(cg)\alpha h > E_{\pi_1}(df)\alpha h.$$

Letting U denote the local MEU utility, we then have: (i) $U(df) > U(cg)$, (ii) $U(cg\alpha h) > U(df\alpha h)$, violating co-monotonic independence.

3 Main Results

The main results of the paper (theorems 1,2 below) are (i) a behavioral characterization of the local MEU model and (ii) an identification theorem for the local MEU model.

Theorem 1. *A preference \succeq satisfies Axioms 1-5 if and only if it admits a local MEU representation.*

The proof is in the appendix. We give a brief sketch of the sufficiency construction as the necessity argument is straightforward.

Sketch of Proof. The proof proceeds in two main steps. First, we show that axioms 1-4 (i.e. just the “standard” axioms, plus strong c -independence, but excluding local uncertainty aversion) imply a representation of the following form. Let $\{\Pi_\alpha\}$ denote a potentially infinite collection of (closed) sets of measures (on $\Delta(\mathcal{S})$) and define:

$$U(f) = \sup_{\alpha \in \Lambda} \text{MEU}_{\Pi_\alpha}(f),$$

where $\text{MEU}_{\Pi_\alpha}(f)$ is the MEU-value of f on the set of measures Π_α and Λ is some index set. This steps is accomplished by looking at the following “hedging” relation. Say that $f \succsim g$ if:

1. $f\alpha g \succeq g$ whenever $f \succeq g$ (and vice-versa, if $g \succeq f$), and
2. whenever $g\alpha h \succeq h$, then $f\alpha h \succeq h$ as well.

The relation \succsim is a well-defined, symmetric¹⁴ binary relation on $B(\mathcal{S})$ (the space of all acts). It is also non-empty, since $f \succsim f$. The relation induces an (undirected) graph on the set $B(\mathcal{S})$ and we can consider the collection \mathcal{K}_α of maximal, complete sub-graphs.¹⁵ Maximal, complete sub-graphs exist by a simple application of Zorn’s lemma. We thus obtain a collection $\{\mathcal{K}_\alpha\}_{\alpha \in \Lambda}$ of maximal, complete sub-graphs of the full graph $(B(\mathcal{S}), \mathcal{E}(B(\mathcal{S})))$, where an undirected edge $\{(f, g), (g, f)\} \in \mathcal{E}(B(\mathcal{S}))$ iff $f \succsim g$. Now fix an f and a maximal component \mathcal{K}_α containing f (which exists since $f \succsim f$, hence lies in some connected component – and thus in some maximal one).¹⁶ Consider the following set:

$$\{d \cdot (u(g) - u(c_f)) : d \geq 0, g \succeq f, g \in \mathcal{K}_\alpha\},$$

where c_f is a certainty equivalent for the act f . From definition of \succsim , and since \mathcal{K}_α is complete we obtain that \mathcal{K}_α is convex. The appendix shows that it is closed as well. Hence, let $\Pi_{\mathcal{K}_\alpha}$ denote a supporting set of (one checks) probability measures. Doing this for each act $f \in \mathcal{K}_\alpha$, taking the union over all such sets of measures (and their closure – call this union $\Pi_{\mathcal{K}_\alpha}$), and finally for each complete (sub)graph \mathcal{K}_α , we then conclude by showing that we recover a representation via the formula

$$(*) U(f) = \sup_{\alpha \in \Lambda} \text{MEU}_{\Pi_{\mathcal{K}_\alpha}}(f).$$

This concludes step 1 of the argument. Importantly, note that no form of uncertainty aversion was used thus far (local or otherwise). To aid the reader, we have

¹⁴The statement given above is an a priori asymmetric, but simpler to state, version of the actual binary relation we use. See the proof of theorem 1 in the appendix for the precise definition of the (symmetric) relation.

¹⁵Recall that an abstract graph is just a pair $(\mathcal{E}, \mathcal{V})$ consisting of a vertex set \mathcal{V} and a set of edges \mathcal{E} connecting (some) of these vertices. A graph is said to be *complete* if any two vertices are connected by an edge between them.

¹⁶Note that, for maximal complete sub-graphs \mathcal{K}_α we can think of these (on the level of sets) as subsets of $B(\mathcal{S})$.

provided – in the appendix – a redone version of the proof of the classical MEU model to fit this outline. In that proof, we also obtain a supporting set of measures attached to the cone supporting each act. As in step 1, absent any other hypotheses there could be a different set of measures supporting each act. Uncertainty aversion is used to show that in fact there is a single supporting set of measures that suffices. In the context of the model (*) uncertainty aversion implies that the collection Λ is a singleton. For our model, we don't have uncertainty aversion – only the weakened version of it. Step 2 of the argument shows that weak uncertainty aversion implies the index set Λ is wlog finite, and this concludes the sketch of the sufficiency argument.

We now turn to identifying the local MEU representation. For this purpose, we label a generic local MEU representation with a pair $(u, \{\Pi_i\}_{i=1}^N)$, where Π_i are (closed, convex) sets of probability measures. We would like to, ideally, say that there is a unique pair $(u, \{\Pi_i\})$ associate to a given preference over acts. Without any further restrictions on the priors, this claim is false. To see this, consider the special case where each Π_i is singleton and consider a three-source local MEU representation, (π_1, π_2, π_3) , where π_3 is a convex combination of π_1, π_2 . Clearly, if we delete the third source we still maintain a representation. This suggests the following restriction on the set of local MEU models.

Definition 4. A local MEU model $(u, \{\Pi_i\})$ is *non-redundant* if there is no sub-collection $\{\Pi'_i\} \subseteq \{\Pi_i\}$ such that the local MEU model $(u, \{\Pi'_i\})$ also represents the same preference.

We will need one more restriction on local MEU models to obtain identification. To explain, consider two representations with source-dependent priors, (π_1, π_2, π_3) and $(\pi_1, \pi_2, \{\pi_3 \cup \{\alpha\pi_1 + (1-\alpha)\pi_2\})$. That is, in the second model we have the same two sources, but in the third source we have two priors, π_3 and $\alpha\pi_1 + (1-\alpha)\pi_2$. However, the value of the MEU functional is never attained on this latter prior (and on this source). Hence, we can delete this prior from the group and maintain the same preference over acts. This suggests the following restriction.

Definition 5. A local MEU model $(u, \{\Pi_i\})$ is *minimal* if there is no other local MEU model $(u, \{\Pi'_i\})$ such that (i) $\Pi_i \supseteq \Pi'_i$ (where the indices match) and (ii) $(u, \{\Pi'_i\})$ represents the same preference over acts.

When a local MEU model is non-redundant and minimal we call it *regular*.

Theorem 2. Let $(u, \{\Pi_i\}_{i=1}^N), (u', \{\Pi'_i\}_{i=1}^M)$ denote two regular local MEU representations of the same preference \succeq . Then, $u = au' + b, a > 0, b \in \mathbf{R}, \text{con}(\Pi) = \text{con}(\Pi')$ and there is a bijection $\zeta : \{\Pi_i\}_{i=1}^N \rightarrow \{\Pi'_i\}_{i=1}^M$ such that $\text{con}(\zeta(\Pi_i)) = \text{con}(\Pi_i)$.

That is, priors are identified up to convexification and the respective collections of subsets of these priors across two representations are in bijection, with members

of the respective collection equal up to convexification. The proof is in the appendix. Here we give a sketch.

Sketch of Proof. The argument proceeds in three main steps, all of which rely on the concept of a *critical act*. Fixing a model, $(u, \{\Pi_i\})$, an act f is critical for this model if the value $U(f)$ is strictly attained on one of the sources Π_i comprising the model. Step 1 of the identification argument shows that is source is determined by the set of acts critical for that source, in the following sense. For each act critical for, say, source Π_1 we can define a set:

$$\{d \cdot (u(g) - u(f)) : d \geq 0, g \succeq f, g \in B_\varepsilon(f)\}.$$

For critical acts f this set turns out to be a closed (convex) cone, for all small ε . Moreover, the cones are ordered by set inclusion as a function of ε . Since taking duals reverses set inclusion, by taking the “limit” as ε goes to zero gives a grand set of measures – which is the limit set of measures supporting the displayed cone. Importantly, since the cone displayed above are defined based only on the primitive, the set of measures supporting the cone (i.e. the dual cone) also depend only on the primitive and not on the particular model $(u, \{\Pi_i\})$ representing this primitive. We then show that (i) for an act f critical for, say, Π_1 its cone of measures is contained in Π_1 and (ii) regularity implies that the Π_1 equals (up to convexification) the closure of the union of these sets of measures, where we take the union over all acts critical for Π_1 . This concludes step 1 of the argument.

Step 2 relies on the identification theorem from [Chandrasekher \(2017\)](#) – which identifies an augmented dual-self model. There is a mathematical connection between the local MEU model and that model since they both exhibit a max-min structure. The point of step 2 is to show that any two regular representations exhibit the same *set* of critical acts. That is, fixing two regular models $(u, \{\Pi_i\}), (u, \{\Pi'_i\})$, the claim is that any act critical for some source Π_i in the first model is also critical for some source Π'_i in the second model, and vice-versa. The way in which the auxiliary identification enters is that we (i) (essentially) take increasingly fine discrete approximations to the space of all acts and (ii) for each discretization, we can reinterpret the local MEU model as a version of the “planner-doer” model from [Chandrasekher \(2017\)](#). The identification result from that paper pins down the collection of critical acts, i.e. they are the same for the common discretization of the two putatively distinct regular representations. By taking increasingly finer discretizations, we obtain equality of critical acts (when we pass to the full space of acts). This turns out to be independent of the sequence of discretizations that we choose at the outset.

Step 3 shows that the set of critical acts determines the model. That is, given step 2, we know that any two regular representations have the same *set* of critical acts. Step 1 also tells us that each source is determined by the set of acts that are

critical for that source. Hence, if we can show that any two acts f, g critical for a common source, say, Π_1 in model $(u, \{\Pi_i\})$ are also critical for a common source in model $(u, \{\Pi'_i\})$, then we obtain an inclusion $\Pi_1 \subseteq \Pi'_{i_1}$. Doing this for each source in the model $(u, \{\Pi_i\})$ gives a source-by-source inclusion of sets of priors. Regularity then implies the sets of measures must agree, concluding the proof of uniqueness.

4 Conclusion

This paper has introduced a new model of ambiguity aversion. The model generalizes the classical max-min expected utility (MEU) model of [Gilboa and Schmeidler \(1989\)](#), with the goal being to explain source-dependent uncertainty aversion. Source-dependence (sometimes also called *source preference*) is the finding that sensitivity to uncertainty is a function of where this uncertainty comes from. There are multiple experimental findings which show that betting preferences exhibit source-dependence. For instance, the well-known competence hypothesis ([Heath and Tversky \(1991\)](#), see also [Fox and Tversky \(1995\)](#)) is based off the finding that decision-makers prefer betting on uncertain but familiar events over those which are less familiar, even when the odds are certain. Our model reconciles uncertainty aversion with source-preference. We show (by example) that our model nests source-preferences as formalized by [Fox and Tversky \(1995\)](#).

5 Appendix

5.1 Preliminaries

We present a preliminary result, an alternative axiomatization of MEU where we replace uncertainty aversion with hedging convexity (defined below).¹⁷ The reason we include this result (which is most certainly already known) is that the proof introduces a method (for eliciting priors) that will be used in all subsequent arguments. It is perhaps useful to first see the construction in the more familiar context of the MEU model.

Axiom 5c: (Hedging Convexity) If $f \sim g, f \sim h$ and $\alpha \cdot f + (1 - \alpha) \cdot g \succeq f, g, \alpha \cdot f + (1 - \alpha) \cdot h \succeq f, h, \forall \alpha \in (0, 1)$, then $\alpha \cdot f + (1 - \alpha) \cdot (\beta \cdot g + (1 - \beta) \cdot h) \succeq f, \forall \alpha, \beta \in (0, 1)$

In words, hedging convexity says that if the DM prefers the hedge of f and g to either one and also a hedge of f and h to either one, then he also prefers a hedge between f and any mixture of g, h to either f or the mixture of g and h . The axiom is formally exactly what the title suggests: the set of acts with which the DM prefers to hedge against a given f is itself a convex set. Hedging convexity weakens uncertainty aversion (UA) but it turns out to be equivalent to the Gilboa-Schmeidler (Gilboa and Schmeidler (1989)) axioms when we combine it with axioms 1-4.

Proposition 2. *A preference \succeq satisfies Axioms 1-5c (i.e. 1,2,3,4, and 5c) if and only if it admits an MEU representation.*

Proof of Proposition 2. Necessity is obvious. Turning to sufficiency, assume the axioms and let $UC(f)$ denote the upper contour set. Under the usual MEU axioms this is a convex and closed set. When we replace uncertainty aversion with hedge convexity, we can still show this is a convex set. This is step 2 below. Temporarily granting the truth of this claim, we derive the representation from this fact.

Step 1: Deriving the set of priors Π assuming convexity of $UC(f)$.

Letting c_f denote (any) certainty equivalent of f we have $UC(f) = UC(c_f)$ and, hence, since both are closed and convex we have $\Pi_f = \Pi_{c_f}$. We claim that $U(f) = \min_{\pi \in \Pi_f} E_{\pi} u(f)$.¹⁸ Consider the cone $\mathcal{C}_{c_f} := \{d \cdot (u(g) - u(c_f)) : d \geq 0, g \succeq f\}$. Since \mathcal{C}_{c_f} is closed (as it clearly has non-empty algebraic interior) it is supported by a (closed, convex) set of probability measures, call it Π_{c_f} , which yields the following characterization:

$$(*) \quad g \in UC(f) \Leftrightarrow E_{\pi} u(g) \geq E_{\pi} u(c_f) = u(c_f), \forall \pi \in \Pi_{c_f}.$$

¹⁷This fact is most likely known to experts in the area, but we have not been able to locate a precise reference, hence, we present the formal statement here.

¹⁸For brevity, we suppress the state in the argument for s , with the understanding that we are really plugging in $f(s)$ into the function $u(\cdot)$.

Let us prove (*). This invokes the following dual characterization of the cone \mathcal{C}_{c_f} :

$$\begin{aligned}\mathcal{C}_{c_f} &= \{d \cdot (u(g) - u(c_f)) : d \geq 0, g \succeq c_f\} \\ &= \{d \cdot (u(g) - u(c_f)) : d \geq 0, E_\pi u(g) \geq E_\pi u(c_f) = c_f, \forall \pi \in \Pi_{c_f}\}.\end{aligned}$$

The equivalence of (*) follows immediately. We claim this implies that $\min_{\pi \in \Pi_{c_f}} E_\pi u(f) = u(c_f)$. Notice that (*) immediately implies the inequality $\min_{\pi \in \Pi_{c_f}} E_\pi u(f) \geq u(c_f)$. We show the reverse inequality via contradiction. Say that $\min_{\pi \in \Pi_{c_f}} E_\pi u(f) > u(c_f)$. Find g such that $\min_{\pi \in \Pi_{c_f}} E_\pi u(g) = u(c_g) > u(c_f)$. Let \underline{c} denote a u -minimal lottery and find α such that $\alpha \cdot u(c_g) + (1 - \alpha) \cdot u(\underline{c}) = u(c_f)$. Consider the act $\hat{f} := \alpha \cdot f + (1 - \alpha) \cdot \underline{c}$. By c -linearity, $f \succ \hat{f}$. OTOH, $\min_{\pi \in \Pi_{c_f}} E_\pi u(\hat{f}) = \min_{\pi \in \Pi_{c_f}} \alpha \cdot E_\pi u(f) + (1 - \alpha) \cdot u(\underline{c}) = \alpha \cdot u(c_g) + (1 - \alpha) \cdot u(\underline{c}) = u(c_f)$. By (*), $\hat{f} \in UC(f)$ – a contradiction.

For the full set of measures, take all constant acts c and consider the collection of priors Π_c . Let $\Pi := \text{co}(\overline{\bigcup_{c \in \Delta(X)} \Pi_c})$, where the closure is taken w.r.t. the weak* topology on the set of measures on \mathcal{S} .¹⁹ Also let $U(\cdot)$ be a cardinal c -linear representation of \succeq and let $u(\cdot)$ denote its restriction to constant acts, i.e. elements of $\Delta(X)$. The existence of U mimics the Gilboa-Schmeidler (Gilboa and Schmeidler (1989)) argument verbatim, hence is omitted. It remains to check that the value $\min_{\pi \in \Pi} E_\pi u(f) = U(f)$ when we look at the expanded set of measures Π . To see this, first fix an f and consider mixing with any constant act c , say $\alpha \cdot f + (1 - \alpha)c$. By c -linearity,

$$U(\alpha \cdot f + (1 - \alpha)c) = \alpha \cdot U(f) + (1 - \alpha)u(c).$$

By the preceding paragraph,

$$\begin{aligned}U(\alpha \cdot f + (1 - \alpha)c) &= \min_{\pi \in \Pi_{\alpha \cdot c_f + (1 - \alpha)c}} E_\pi u(\alpha \cdot f + (1 - \alpha)c) \\ &= \min_{\pi \in \Pi_{\alpha \cdot c_f + (1 - \alpha)c}} \alpha \cdot E_\pi u(f) + (1 - \alpha)u(c).\end{aligned}$$

Hence, canceling common terms, we obtain: (**) $U(f) = \min_{\pi \in \Pi_{\alpha \cdot c_f + (1 - \alpha)c}} E_\pi u(f)$. On the other hand, we have $U(f) = \min_{\pi \in \Pi_{c_f}} E_\pi u(f)$. By (**), we have

$$\min_{\pi \in \Pi_{\alpha \cdot c_f + (1 - \alpha)c}} E_\pi u(f) = \min_{\pi \in \Pi_{c_f}} E_\pi u(f).$$

Since c is an arbitrary element of $\Delta(X)$, we obtain $\min_{\pi \in \Pi} E_\pi u(f) = U(f)$.

Step 2: Convexity of $UC(f)$.

We show indirectly that $UC(f)$ is convex (when hedging convexity replaces uncertainty aversion). Let $UC^0(f)$ denote the subset of $UC(f)$ for which hedging with f

¹⁹*N.B.* This is a finite set, hence this is the same as the Euclidean topology.

has value. By axiom 5a and continuity (axiom 2) we know that $UC^0(f)$ is convex (and closed). Now for any f let c_f denote a certainty equivalent for f and observe that

$$UC^0(c_f) = UC(c_f) = UC(f).$$

The last equality follows by definition of c_f . For the first equality, take any $g \in UC(c_f)$ and note that c -independence implies $\alpha \cdot g + (1 - \alpha) \cdot c_f \succeq \alpha \cdot c_f + (1 - \alpha) \cdot c_f = c_f, \forall \alpha \in [0, 1]$. Hence, $g \in UC^0(c_f)$. Since $UC^0(c_f)$ is convex (for any act, not just the constants) it follows that $UC(f)$ is convex as well (and closed). \square

The key step in this proof was the convexity of $UC(f)$. This allowed us to show that the value of the c -linear functional $U(\cdot)$ at f was recovered by minimizing expected utility on the set of measures supporting the upper contour set of f . When we replace hedge convexity (or uncertainty aversion) with weak uncertainty aversion, the upper contour set is *not* convex. Nevertheless, we will be able to recover the equality $U(f) = \min_{\pi \in \Pi_f} E_\pi u(f)$, where the measures comprising Π_f support a modified cone. To define the modification put

$$UC^0(f) := \{g \in UC(f) : \alpha \cdot g + (1 - \alpha) \cdot f \succeq f, \forall \alpha\}.$$

That is, $UC^0(f)$ consists of all acts weakly preferred to f which provide hedging value against f . We will actually look at a set $UC^0(f)$ which not only provides hedging value against f (and which is preferred to f) but the subset of this that satisfies hedging convexity. For the precise expression of this subset we defer to the actual proof of theorem 1.

5.2 Proofs for section 2

No MEU representation for example 2. Let $(\underline{x}, \bar{x}, \underline{y}, \bar{y}, \underline{z}, \bar{z})$ denote a putative solution to the system of inequalities. We normalize the inequalities to embed the solution space of the inequality system into the solution space of an LP problem. Coefficients are bolded for ease of reference. We use the symmetry of the inequalities to break up the coefficient matrix into three pieces. There will be, all told, 24 equalities that are to be satisfied together by the sextuplet $\{\underline{x}, \bar{x}, \underline{y}, \bar{y}, \underline{z}, \bar{z}\}$, along with s_1, \dots, s_{24} surplus (dummy) variables.

$$\begin{aligned} \mathbf{1} \cdot \underline{x} + \mathbf{0} \cdot \bar{x} - \mathbf{1} \cdot \underline{z} + \mathbf{0} \cdot \bar{z} - \mathbf{1} \cdot s_1 + \dots + \mathbf{0} \cdot s_7 &= 0 \\ \mathbf{0} \cdot \underline{x} + \mathbf{1} \cdot \bar{x} + \mathbf{0} \cdot \underline{z} - \mathbf{1} \cdot \bar{z} + \mathbf{1} \cdot s_2 + \dots + \mathbf{0} \cdot s_7 &= 0 \\ \mathbf{0} \cdot \underline{x} + \mathbf{1} \cdot \bar{x} + \mathbf{1} \cdot \underline{z} + \mathbf{1} \cdot s_3 + \dots + \mathbf{0} \cdot s_7 &= 1 \\ \mathbf{1} \cdot \underline{x} + \mathbf{0} \cdot \bar{x} + \mathbf{1} \cdot \bar{z} + \mathbf{0} \cdot s_1 + \dots + \mathbf{1} \cdot s_4 + \mathbf{0} \cdot s_7 &= 1 \\ \mathbf{1} \cdot \underline{x} - \mathbf{1} \cdot \bar{x} + \mathbf{1} \cdot s_5 + \mathbf{0} \cdot s_7 &= 0 \\ \mathbf{0} \cdot \underline{x} + \mathbf{1} \cdot \bar{x} + \mathbf{1} \cdot \bar{z} - \mathbf{2} \cdot \underline{z} + \mathbf{0} \cdot s_1 \dots - \mathbf{1} \cdot s_6 + \mathbf{0} \cdot s_7 &= 0 \\ \mathbf{0} \cdot \underline{x} + \mathbf{0} \cdot \bar{x} + \mathbf{1} \cdot \underline{z} - \mathbf{1} \cdot \bar{z} + \mathbf{1} \cdot s_7 &= 0. \end{aligned}$$

Equations (5), (12) and (7), (21) and (14),(19) are replicas. Replicating equations doesn't impact (non)solubility of an LP. They are both included just to obtain a symmetric system, e.g. when we swap out x for y in the third system below. Now consider the (symmetric) set of equations involving the variables $(\underline{x}, \bar{x}, \underline{y}, \bar{y})$. This gives, labelling the dummy variables for this block, s_8, \dots, s_{14} .

$$\begin{aligned}
1 \cdot \underline{x} + 0 \cdot \bar{x} - 1 \cdot \underline{y} + 0 \cdot \bar{y} - 1 \cdot s_8 + \dots + 0 \cdot s_{14} &= 0 \\
0 \cdot \underline{x} + 1 \cdot \bar{x} + 0 \cdot \underline{y} - 1 \cdot \bar{y} + 1 \cdot s_9 + \dots + 0 \cdot s_{14} &= 0 \\
0 \cdot \underline{x} + 1 \cdot \bar{x} + 1 \cdot \underline{y} + 1 \cdot s_{10} + \dots + 0 \cdot s_{14} &= 1 \\
1 \cdot \underline{x} + 0 \cdot \bar{x} + 1 \cdot \bar{y} + 0 \cdot s_8 + \dots + 1 \cdot s_{11} + 0 \cdot s_{14} &= 1 \\
1 \cdot \underline{x} - 1 \cdot \bar{x} + 1 \cdot s_{12} + 0 \cdot s_{14} &= 0 \\
0 \cdot \underline{x} + 1 \cdot \bar{x} + 1 \cdot \bar{y} - 2 \cdot \underline{y} + 0 \cdot s_8 \dots - 1 \cdot s_{13} + 0 \cdot s_{14} &= 0 \\
0 \cdot \underline{x} + 0 \cdot \bar{x} + 1 \cdot \underline{y} - 1 \cdot \bar{y} + 1 \cdot s_{14} &= 0.
\end{aligned}$$

Now consider the set of equations involving $(\underline{y}, \bar{y}, \underline{z}, \bar{z})$.

$$\begin{aligned}
1 \cdot \underline{y} + 0 \cdot \bar{y} - 1 \cdot \underline{z} + 0 \cdot \bar{z} - 1 \cdot s_{15} + \dots + 0 \cdot s_{21} &= 0 \\
0 \cdot \underline{y} + 1 \cdot \bar{y} + 0 \cdot \underline{z} - 1 \cdot \bar{z} + 1 \cdot s_{16} + \dots + 0 \cdot s_{21} &= 0 \\
0 \cdot \underline{y} + 1 \cdot \bar{y} + 1 \cdot \underline{z} + 1 \cdot s_{17} + \dots + 0 \cdot s_{21} &= 1 \\
1 \cdot \underline{y} + 0 \cdot \bar{y} + 1 \cdot \bar{z} + 0 \cdot s_{15} + \dots + 1 \cdot s_{18} + 0 \cdot s_{21} &= 1 \\
1 \cdot \underline{y} - 1 \cdot \bar{y} + 1 \cdot s_{19} + 0 \cdot s_{21} &= 0 \\
0 \cdot \underline{y} + 1 \cdot \bar{y} + 1 \cdot \bar{z} - 2 \cdot \underline{z} + 0 \cdot s_{15} \dots - 1 \cdot s_{20} + 0 \cdot s_{21} &= 0 \\
0 \cdot \underline{y} + 0 \cdot \bar{y} + 1 \cdot \underline{z} - 1 \cdot \bar{z} + 1 \cdot s_{21} &= 0.
\end{aligned}$$

Finally, we have the following three inequalities involving all six of the variables $(\underline{x}, \bar{x}, \dots, \underline{z}, \bar{z})$:

1. $\underline{x} > 1 - (\bar{y} + \bar{z})$.
2. $\underline{z} > 1 - (\bar{x} + \bar{y})$.
3. $\underline{y} > 1 - (\bar{x} + \bar{z})$.

Standardizing gives the following matrix form:

$$\begin{aligned}
1 \cdot \underline{x} + 1 \cdot \bar{y} + 1 \cdot \bar{z} - 1 \cdot s_{22} &= 1 \\
1 \cdot \bar{x} + 1 \cdot \bar{y} + 1 \cdot \underline{z} - 1 \cdot s_{23} &= 1 \\
1 \cdot \bar{x} + 1 \cdot \underline{y} + 1 \cdot \bar{z} - 1 \cdot s_{24} &= 1.
\end{aligned}$$

Let A denote the coefficient matrix of the entire system and break the rows of A into A_1, A_2, A_3, A_4 where $A_1 - A_3$ are each 7×11 and A_4 is 3×9 (for A_i we add zeroes

in the columns for dummy variables involved only in $A_j, j \neq i$). Non-representability by the MEU model will be implied if we show the alternative to the preceding LP (whose solution space contains the solutions to the MEU inequalities) is non-empty. To show non-emptiness of the alternative we compute the two matrix products, yA and $b \cdot y$, where $b = [0, 0, 1, 1, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 1, 1, 1]$. First we compute the column products for the dummy variables:

1. $y_2 - y_5, y_7, y_9 - y_{12}, y_{14}, y_{16} - y_{19}, y_{21} \leq 0$,
2. $y_1, y_6, y_8, y_{13}, y_{15}, y_{20}, y_{22} - y_{24} \geq 0$,

Now compute the product $b \cdot y$: (negative numbers are bolded)

$$(*) \mathbf{y_3} + \mathbf{y_4} + \mathbf{y_{10}} + \mathbf{y_{11}} + \mathbf{y_{17}} + \mathbf{y_{18}} + y_{22} + y_{23} + y_{24} > 0.$$

Now list the matrix product of y with column i in order.

1. $y_1 + \mathbf{y_4} + \mathbf{y_5} + \mathbf{y_8} + \mathbf{y_{11}} + y_{12} + y_{22} \leq 0$.
2. $\mathbf{y_2} + \mathbf{y_3} + \mathbf{y_5} + y_6 + \mathbf{y_9} + \mathbf{y_{10}} + \mathbf{y_{12}} + y_{13} + y_{23} + y_{24} \leq 0$.
3. $y_8 + \mathbf{y_{10}} - 2y_{13} + \mathbf{y_{14}} + y_{15} + \mathbf{y_{18}} + \mathbf{y_{19}} + y_{24} \leq 0$.
4. $-y_9 + \mathbf{y_{11}} + y_{13} - y_{14} + \mathbf{y_{16}} + \mathbf{y_{17}} - y_{19} + y_{20} + y_{22} + y_{23} \leq 0$.
5. $-y_1 + \mathbf{y_3} + \mathbf{y_7} - 2y_6 - y_{15} + \mathbf{y_{17}} - 2y_{20} + y_{21} + y_{23} \leq 0$.
6. $-y_2 + \mathbf{y_4} + y_6 + \mathbf{y_7} - y_{16} + \mathbf{y_{18}} + y_{20} - y_{21} + y_{22} + y_{24} \leq 0$.

To show that $(*)$ and each of (1) – (6) can be satisfied (in a manner consistent with the sign conditions, it suffices to show that for each of $(*)$ and (1) – (6) we can choose a dominant term, i.e the sign of the corresponding sum is determined by the sign of a putative dominant term. Dominant terms are not necessarily unique, but we just need to obtain one selection of such terms across $(*)$ and (1) – (6). The constraint on the selection is that it must be internally consistent, e.g. if we select y_j to be dominant in, say, equation (6) where y_i is also present, then we cannot have y_i be dominant in equation (5) where both y_i, y_j are still present. The following is a (consistent) selection of dominant terms:

- For $(*)$ choose y_{23} as dominant.
- For (6), y_4 is dominant.
- For (5), y_6 is dominant.
- For (4), y_{16} is dominant.

- For (3), y_{13} is dominant.
- For (2), y_2 is dominant.
- For (1), y_5 is dominant.

It follows that the alternative to the LP program is non-empty, implying there is no MEU representation of the source preference. \square

Proof of Proposition 1. Denote θ -local MEU representation with the triple $(u, \{\Pi_i\}, \theta)$ and local MEU representations via $(u, \{\Pi_i\})$. Obviously finite local MEU representations are θ -MEU representations, so we check that the former is also nested in the latter when we allow the local MEU model to have potentially infinite sources. That is, given $(u, \{\Pi_i\}, \theta)$ we find a tuple $(u, \{\Pi'_i\})$ such that the two models yield the same cardinal utility over acts. Proceed in two steps. First, for each act f , we find a subdivision of the initial sources $\{\Pi_i\}$, where each set Π_i is divided into several (possibly overlapping) $\{\Pi_i^j(f)\}$, where $\cup_j \Pi_i^j(f) = \Pi_i$. Second, we show how these various subdivisions can be further amended and fit together to obtain a single subdivision for all acts.

Step 1: Defining a single subdivision, $\{\Pi_i^j(f)\}_i$.

For each act f , break up the collection $\{\Pi_i\}$ into two sets, $\Sigma_1(f), \Sigma_2(f)$ according to whether the θ -MEU value $U(f)$ is attained on the set Π_i or not. Let $\Sigma_2(f)$ denote the sub-collection of those Π_i on which the value is not attained. Let $\Sigma_1(f)$ denote those Π_i on which the θ -MEU value of $U(f)$ is attained. Fix an $\Pi_i \in \Sigma_1(f)$ and replace with the following collection. Let v_θ be a numerical representation of θ , i.e. $\theta(A) = \arg \max_{x \in A} v_\theta(x)$ for all $A \subseteq \mathbf{R}$ that are compact, and consider for each $z \in \{E_\pi u(f) : \pi \in \Pi_i\}$, the following set:

$$\Pi_i^z(f) = \{\pi \in \Pi_i : E_\pi u(f) \leq z, v_\theta(E_\pi u(f)) \geq v_\theta(z)\}$$

This gives a (non-partitional) subdivision of Π_i into $\{\Pi_i^z(f)\}_z$. We now claim that

$$(*) \max_z \min_{\pi \in \Pi_i^z(f)} E_\pi u(f) = \theta - \text{MEU}_{\Pi_i}(f).$$

In other words, the θ -local MEU value of f is attained by, alternatively, subdividing the set of measures Π_i as given above and apply the MEU functional to this collection recovers the θ -MEU value. Let us check this. For each set $\Pi_i^z(f)$ we an associated set of expected values, $\{E_\pi u(f) : \pi \in \Pi_i^z(f)\}$. Note that

$$\theta(\underbrace{\{E_\pi u(f) : \pi \in \Pi_i^z(f)\}}_{A_z^i(f)}) = \arg \max_{x \in \{E_\pi u(f) : \pi \in \Pi_i^z(f)\}} v_\theta(x).$$

Also observe that, for $z = \theta - \text{MEU}(f)$, we have $\theta(A_z^i(f)) = \theta - \text{MEU}(f)$, by property (iii) defining the map θ (which is what gives the numerical representation v_θ). Since, on the set $\Pi_i^z(f)$, we have $\min_{\pi \in \Pi_i^z(f)} E_\pi u(f) = \theta(A_z^i(f))$ it follows

that the MEU value of f on $\Pi_z^i(f)$ equals its θ -MEU value. Since we are making further θ -selections from and then taking the maximum across sets $A_z^i(f)$, we clearly obtain $\max_z \min_{\pi \in \Pi_z^i(f)} E_\pi u(f) \geq \theta - \text{MEU}_{\Pi_i}(f)$. For the reverse inequality, towards a contradiction say that there is some set $\Pi_z^i(f)$, with associated $A_z^i(f)$ for which $\theta(A_z^i(f)) > \theta - \text{MEU}(f)$. Let $\pi_f \in \Pi_z^i(f)$ denote some measure on which the $\theta - \text{MEU}$ value is attained. Note that $E_{\pi_f} u(f) \leq z$. Moreover, since $A_z^i(f) = \{E_\pi u(f) : \pi \in \Pi_z^i(f)\} \subseteq \{E_\pi u(f) : \pi \in \Pi_i\}$ and v_θ is a numerical representation of the selection map $\theta(\cdot)$, it follows that $\theta(A_z^i(f)) = \theta - \text{MEU}(f)$. This proves the equality (*).

Step 2: Finding a single coherent subdivision, $\{\Pi_i^j\}_j$, for all acts.

For each act g we have a subdivision $\{\Pi_i^j(g)\}$ of the original collection $\{\Pi_i\}$ as described in step 1, i.e. where each $\Pi_i^j(g) = \{\pi \in \Pi_i : E_\pi u(g) \leq z\}$ for some z (and where, wlog, we restrict to z in the range of $E \cdot u(g)$). For each act f , consider sets $\Pi_i(f)$ on which the $\theta - \text{MEU}$ value of f is attained. Denote these sets as $\Pi_i^*(f)$. Fix a $\Pi_i^*(f)$. For each g with $\theta - \text{MEU}(g) > \min_{\pi \in \Pi_i^*(f)} E_\pi u(g)$ we claim that we can find a measure $\pi(f, g)$ such that:

1. $E_{\pi(f, g)} u(f) \geq \theta - \text{MEU}(f)$, and
2. $E_{\pi(f, g)} u(g) < \theta - \text{MEU}(g)$.

It is straightforward to check that such a measure $\pi(f, g)$ exists. To see this, let $u_\theta(f), u_\theta(g)$ denote the (resp.) local θ -MEU values of f, g . Note that, by replacing f with $f \alpha c$ we can assume $u_\theta(f) = u_\theta(g) (= \beta)$. Note that, since $f \neq g$, this implies $f \neq \lambda \cdot g$. Otherwise, $f = \lambda \cdot g$. Hence, the θ -local MEU value occurs on the same measure π (since $\lambda > 0$) and the equality of these values forces $\lambda = 1$ – a contradiction. Now we claim the following system has a positive solution, $\{\pi_s\}$, for small ε :

$$\begin{aligned} \sum_s \pi_s \cdot u(f(s)) &= \beta + \varepsilon \\ \sum_s \pi_s \cdot u(g(s)) &= \beta - \varepsilon. \end{aligned}$$

Since $f \neq \lambda \cdot g$, it is clear that the system has a solution for some $\{\pi_s\}$. We check that there is, additionally, a positive solution. Via Farkas' lemma, if no positive solution exists, then there is some $y = (y_1, y_2)$ such that, letting $A = [u(f), u(g)]^t, b = [\beta + \varepsilon, \beta - \varepsilon]^t, y \cdot A \leq 0, b \cdot y > 0$. Proceeding towards contradiction, consider two cases, (i) $\exists (s_1, s_2), u(f(s_1)) \geq u(g(s_1)), u(f(s_2)) \leq u(g(s_2))$ and (ii) $\nexists (s_1, s_2), u(f(s_1)) \geq u(g(s_1)), u(f(s_2)) \leq u(g(s_2))$, i.e. in the latter case, either f state-by-state dominates g or vice-versa. As we shrink ε , consider candidate elements y_ε of the Farkas alternative. Normalize such that $\|y_\varepsilon\| = 1$ and pass, if necessary, to a convergent subsequence. Since $b_\varepsilon \cdot y_\varepsilon > 0$, for all small ε we must have either (i) $|y_1| \geq |y_2|, y_1 \geq 0$, or

(ii) $|y_2| \geq |y_1|$, $y_2 \geq 0$ (for all large ε). Say the former holds and consider a state s_1 in which $u(f(s_1)) \geq u(g(s_1))$. Since $yA \leq 0$, we must have $y_1 u(f(s_1)) + y_2 u(g(s_1)) \leq 0$ – contradiction. If (ii) holds, consider a state s_2 such that $u(f(s_2)) \leq u(g(s_2))$. Since $yA \leq 0$, we must have $y_1 u(f(s_2)) + y_2 u(g(s_2)) \leq 0$ – again a contradiction. Hence, the Farkas alternatives must be empty for all ε small – implying a positive solution to the above system. Let $\{\pi'_s\}$ denote a positive solution and let $\pi_s := \pi'_s / \sum_s \pi'_s$ denote its normalization. This yields the desired measure $\pi(f, g)$, in the case where neither f nor g state-by-state dominates the other. Say that g state-by-state dominates f . Note that (from step 1) we found that $\theta - \text{MEU}_{\Pi_i^*(f)}(f) = \text{MEU}_{\Pi_i^*(f)}(f)$. If additionally $u_\theta(f) = u_\theta(g)$, then we must have $u_\theta(g) = \text{MEU}_{\Pi_i^*(f)}(f)$. On the other hand, by the fact that $\theta(A) \in A$, we must have, $u_\theta(g) \geq \theta - \text{MEU}_{\Pi_i^*(f)}(f) \geq \text{MEU}_{\Pi_i^*(f)}(g)$ – where the latter inequality is strict by hypothesis. Put together we get, $\text{MEU}_{\Pi_i^*(f)}(f) > \text{MEU}_{\Pi_i^*(f)}(g)$ – contradicting the fact that g state-by-state dominates f . Thus, we obtain $\pi(f, g)$ with the requisite properties. Replace $\Pi_i^*(f)$ with the set

$$\Pi_i^{**}(f) = \Pi_i^*(f) \cup \bigcup_{\{g: \theta - \text{MEU}(g) < \min_{\pi \in \Pi_i^*(f)} E_\pi u(g)\}} \{\pi(f, g)\}$$

As before, we then take convexification and closure but suppress this and refer to the resulting set of measures as $\Pi_i^{**}(f)$. Having done this for each f , replace the collection of sub-divisions $\{\Pi_i^j(f)\}_f$ with the single collection, $\{\Pi_i^{**}(f)\}_f$, i.e. if there are multiple i 's on which f attains its value we simply choose one such i and omit all others. By construction, $u_\theta(f) = \text{MEU}_{\Pi_i^{**}(f)}(f), \forall f$ and $\text{MEU}_{\Pi_i^{**}(f)}(g) < u_\theta(g)$ for all $g \neq f$ such that $\theta - \text{MEU}_{\Pi_i^*(f)}(g) < \text{MEU}_{\Pi_i^*(f)}$. It follows that the local MEU model formed from the collection $(u, \{\Pi_i^{**}(f)\}_f)$ generates the same cardinal utility on acts, viz. $u_\theta(f), \forall f$. \square

5.3 Proofs for section 3

Proof of Theorem 1. Necessity is only non-trivial for weak uncertainty aversion. To check this, fix any f and let $\kappa_{i_1}(\Pi), \dots, \kappa_{i_k}(\Pi)$ denote the sources on which the value $U(f)$ is attained. Put $U_j(f) := \min_{\pi \in \kappa_{i_j}(\Pi)} E_\pi u(f)$ and note that

$$U_j(f) > \max_{\kappa_{i_i}(\Pi) \notin \{\kappa_{i_1}(\Pi), \dots, \kappa_{i_k}(\Pi)\}} E_\pi u(f).$$

Find ϵ_f small enough such that $U_j(f) - \epsilon > \max_{\kappa_{i_i}(\Pi) \notin \{\kappa_{i_1}(\Pi), \dots, \kappa_{i_k}(\Pi)\}} E_\pi u(f)$. Consider $B_{\epsilon_f}(f)$ and note that for any $g \in B_{\epsilon_f}(f)$ the value $U(f)$ is attained on one of the sets $\{\kappa_{i_1}(\Pi), \dots, \kappa_{i_k}(\Pi)\}$. It follows that, if $g \sim f$, then on some $\kappa_{i_j}(\Pi)$ at which the value $U(f), U(g)$ is attained we have $U_j(\alpha \cdot f + (1 - \alpha) \cdot g) \geq U(f) = U(g), \forall \alpha \in [0, 1]$. Hence, $U(\alpha \cdot f + (1 - \alpha) \cdot g)$ (the maximum value across all sources) exceeds $U(f), U(g)$.

Now we turn to sufficiency. Introduce $f \alpha g := \alpha \cdot f + (1 - \alpha) \cdot g$. Define a binary relation, \succsim , as follows. Say that $f \succsim g$ if the following two properties hold:

1. If $f \succeq g$, then for any c such that $f_c = \alpha \cdot f + (1 - \alpha) \cdot c \sim g$ we have $\beta \cdot f_c + (1 - \beta) \cdot g \succeq g, \forall \beta \in (0, 1]$.
2. If $\exists h$ such that $g \succeq h$ and $\exists c$ such that $g_c = \alpha \cdot g + (1 - \alpha) \cdot c \sim h$ and $\beta \cdot g_c + (1 - \beta) \cdot h \succeq g_c, h, \forall \beta$, then $\forall \alpha$ we have $\beta \cdot f \alpha g_c + (1 - \beta) h \succeq h, (f \alpha g)_c, \forall \beta$.²⁰

Think of the full binary relation \succsim as a directed graph with vertices taken to be elements of $B(\mathcal{S})$. Connect two vertices, f, g , if and only if $f \succsim g$. This defines a set of (undirected) edges, $\mathcal{E}(B(\mathcal{S}))$. Denote the resulting graph, \mathcal{G}_{\succsim} , with the pair $(B(\mathcal{S}), \mathcal{E}(B(\mathcal{S})))$. Let $\Sigma(B(\mathcal{S}), \mathcal{E}(B(\mathcal{S})))$ denote the set of *complete* sub-graphs of $(B(\mathcal{S}), \mathcal{E}(B(\mathcal{S})))$. A standard application (omitted) of Zorn's lemma shows that this (non-empty) set possesses maximal elements. Let $\{\mathcal{K}_\alpha\}_{\alpha \in \Lambda}$ denote the collection maximal complete sub-graphs. We now proceed in two steps. First, we associate a cone of measures to each \mathcal{K}_α . Second, we show that the set Λ is finite – i.e. there are only finitely many maximal subsets of the graph \mathcal{G}_{\succsim} . It is this fact, and this fact alone, that uses the additional structure on the preference imposed by weak uncertainty aversion.

Step 1: Recovering a cone of measures supporting \mathcal{K}_α .

For each \mathcal{K}_α and each $f \in B(\mathcal{S})$ define the following set:

$$\mathcal{C}_f(\mathcal{K}_\alpha) := \{d \cdot (u(g) - u(f)) : d \geq 0, g \succeq f, (g, f) \in \mathcal{K}_\alpha\}.$$

By construction of the graph, the set $\mathcal{C}_f(\mathcal{K}_\alpha)$ must be convex. We verify the cone must be closed as well. To show this note that the condition $f \succsim g$ is equivalent to the following condition. Say $f \succeq g$ and take $U(\cdot)$ to be a c -linear, monotonous, homogenous cardinal representation of \succeq . Then, $f \succsim g$ if:

1. $U(f \alpha g) \geq \min\{U(f), U(g)\}$, and
2. $U(f \alpha h) \geq \min\{U(f), U(h)\}$

whenever $f \succeq g$ and $U(g \alpha h) \geq \min\{U(g), U(h)\}$. Now take any $g \in \mathcal{K}_\alpha$ for which $g \succ f$. By mixing with a u -maximal constant we can guarantee that $g \succ f$. For small ε consider the vector $\varepsilon \mathbf{1}_{e_i}$ (where $e_1, \dots, e_k, k = |\mathcal{S}|$ are the coordinate bases). The perturbed utility vector $u(g) + \varepsilon \cdot \mathbf{1}_{e_i}$ is itself a utility vector for small ε , denoted g_{ε_i} for short. Note that $U(g_{\varepsilon_i} \alpha f) \geq U(g \alpha f) \geq U(g)$ (the former inequality by monotonicity and the latter since $g \succ f$ and $g \succeq f$). Moreover, whenever $f \succeq h$ and $U(f \alpha h) \geq \min\{U(f), U(h)\}$ we have (by hypothesis) $U(g \alpha h) \geq \min\{U(g), U(h)\} = U(h)$. By monotonicity, we obtain $U(g_{\varepsilon_i} \alpha h) \geq U(g \alpha h) \geq U(h) = \min\{U(g_{\varepsilon_i}), U(h)\}$. Hence, $g_{\varepsilon_i} \succ f$ and $g_{\varepsilon_i} \succeq f$. Hence, the linear span of the (convex) set $\{g : g \succeq f, g \in \mathcal{K}_\alpha\}$ has full dimension (in \mathbf{R}^k), implying that the cone $\{d \cdot (u(g) - u(f)) : d \geq 0, g \in$

²⁰The notation $(f \alpha g)_c$ denotes the mixture of $f \alpha g$ which is indifferent to h .

$\mathcal{K}_\alpha, g \succeq f\}$ is closed. The same argument verbatim also shows that the (closed) set of measures supporting the cone must be (after normalization) probability measures. Let $\Pi_f(\mathcal{K}_\alpha)$ denote the measures supporting the cone.

We now invoke the following (dual) characterization of $\mathcal{C}_f(\mathcal{K}_\alpha)$:

$$\begin{aligned} \{d \cdot (u(g) - u(f)) : d \geq 0, g \succeq f, g \in \mathcal{K}_\alpha\} &= \\ &= \{d \cdot (u(g) - u(f)) : d \geq 0, E_\pi u(g) \geq E_\pi u(f), \forall \pi \in \Pi_f(\mathcal{K}_\alpha)\}. \end{aligned}$$

Let c_f denote a certainty equivalent for f and, for each \mathcal{K}_α with $f \in \mathcal{K}_\alpha$ ²¹ consider the set $\Pi_{c_f}(\mathcal{K}_\alpha)$. Applying the dual characterization we get $\min_{\pi \in \Pi_{c_f}(\mathcal{K}_\alpha)} E_\pi u(f) \geq u(c_f)$. We claim this holds with equality. Else, let $\min_{\pi \in \Pi_{c_f}(\mathcal{K}_\alpha)} E_\pi u(f) = u(\hat{c}) > u(c_f)$ and find c, α with $\alpha \cdot u(\hat{c}) + (1 - \alpha) \cdot c = u(c_f)$. Put $\hat{f} := \alpha \cdot f + (1 - \alpha) \cdot c$ and note that, since $u(c) < u(c_f)$ we have (by c -linearity) $\hat{f} \prec c_f$. On the other hand, we have $\min_{\pi \in \Pi_{c_f}(\mathcal{K}_\alpha)} E_\pi u(\hat{f}) = \min_{\pi \in \Pi_{c_f}(\mathcal{K}_\alpha)} \alpha \cdot E_\pi u(f) + (1 - \alpha) \cdot u(c) = \alpha \cdot u(\hat{c}) + (1 - \alpha) \cdot u(c) = u(c_f)$. Hence, by the dual characterization of $\mathcal{C}_{c_f}(\mathcal{K}_\alpha)$ we have $\hat{f} \succeq c_f \sim f$ – a contradiction, since $\hat{f} \prec c_f$ (and $\hat{f} \in \mathcal{K}_\alpha$). Let

$$\Pi := \overline{\cup_{\alpha \in \Lambda} \cup_{c_f, f \in \mathcal{K}_\alpha} \Pi_{c_f}(\mathcal{K}_\alpha)},$$

where (as before) the closure is taken w.r.t. the weak* topology on the set of measures on \mathcal{S} (which is just Euclidean convergence since the state space is finite). Note that the same argument (i.e. the dual characterization of the cones $\mathcal{C}_c(\mathcal{K}_\alpha)$) shows that whenever $f \succeq c$ and $f \notin \mathcal{K}_\beta$, then we must have $\min_{\pi \in \Pi_c(\mathcal{K}_\beta)} E_\pi u(f) < u(c)$. Hence, we obtain a representation of U via $U(f) = \sup_{\alpha \in \Lambda} \min_{\pi \in \Pi_c(\mathcal{K}_\alpha)} E_\pi u(f)$. We now turn to showing the index set Λ is finite.

Step 2: Checking that the index set Λ is (wlog) finite.

Towards a contradiction, say that there is an infinite collection $\Pi_c(\mathcal{K}_\alpha)$ with the property that the MEU-values of f along this collection increase to $U(f)$.²² To each f_n add a constant act ε_n so that $U(f_n + \varepsilon_n) = U(f)$. Filter the collection of pivotal sets as follows. Pick any sequence $\beta_n \downarrow 0$ and consider the sets

$$\Lambda^0(\beta_n) := \{\alpha \in \Lambda : \exists f, \text{MEU}_{\Pi_c(\mathcal{K}_\alpha)}(f) \leq \text{MEU}_{\Pi_c(\mathcal{K}_{\alpha'})}(f) - \beta_n, \forall \alpha' \in \Lambda\}.$$

We show Λ is finite in a sequence of sub-steps.

²¹By this we mean that there is some pair $(f, h) \in \mathcal{K}_\alpha$ (equivalently, (h, f) since the graph is undirected) with $f \succ h$.

²²The latter claim is wlog. Clearly we can pass to a convergent subsequence. Fixing any convergent subsequence, we can either find a subsequence decreasing to $U(f)$ or one increasing to $U(f)$. It may be the case that we can only find convergent subsequences decreasing to $U(f)$. The argument in this case is entirely symmetric to the one presented, hence omitted.

Sub-Step 2a. Showing that $f \in \Lambda^0(\beta_n)$ for some n .

Towards contradiction, say f is such that along a sequence of sets $\Pi_c(\mathcal{K}_n)$ we have $\text{MEU}_{\Pi_c(\mathcal{K}_n)}(f) \uparrow U(f)$. Let $\pi_n \in \Pi_c(\mathcal{K}_n)$ be an associated sequence of measures on which the MEU-values are attained. Fixing any N we must have $\pi_m \notin \overline{\text{co}(\cup_{n \geq N} \pi_n)}$, $\forall m \leq N$ (since the values $\text{MEU}_{\Pi_c(\mathcal{K}_n)}(f)$ are increasing). For the given f , find ε_f such that, via axiom 5, f' has hedging value with f whenever $f' \in B_{\varepsilon_f}(f)$. Let N be large enough so that $U(f) - U(f_n) < \varepsilon_f/3$, $\forall n \geq N$. Consider the closed convex sets, $\Pi_1^N := \overline{\text{co}(\cup_{n \geq N+1} \pi_n)}$, $\Pi_2^N := \Pi_c(\mathcal{K}_N)$. Consider two cases,

- i. $\exists N, \Pi_1^N \cap \Pi_2^N = \emptyset$, and
- ii. $\forall N, \Pi_1^N \cap \Pi_2^N \neq \emptyset$.

Case (i):

First consider case (i). By the separating hyperplane theorem, there is a vector v_N such that $v_N \cdot \pi^1 < 0 < v_N \cdot \pi^2$, $\forall \pi^1 \in \Pi_1, \forall \pi^2 \in \Pi_2$ (we suppress the N for which the sets are disjoint here). Note that we have used the fact that the sets being separated are hulls of probability measures to normalize the separation constant to be 0. We further normalize so that $\|v_N\| = 1$. Consider the two acts, $\{f, f_N\}$, where $f_N = f + \varepsilon_N \cdot v_N$, $\varepsilon_N = U(f) - U(f_N) + \varepsilon_f/3$. Note that $f_N \in B_{\varepsilon_f}(f)$. Note that the MEU-value of f_N on the sets $\Pi_c(\mathcal{K}_n)$, $\forall n \geq N + 1$ is bounded above by $E_{\pi_n} u(f_N) = E_{\pi_n} u(f) + \varepsilon_N \pi_n \cdot v_N < E_{\pi_n} u(f) < U(f)$. Also note that the MEU value of f_N on $\Pi_c(\mathcal{K}_N)$ exceeds $U(f)$. Hence, by axiom 5, $f \alpha f_N \succ f$, $\forall \alpha \in (0, 1)$. However, find α large enough so that the MEU-value of $f \alpha f_N$ on $\Pi_c(\mathcal{K}_m)$, $\forall m \leq N$ is less than $U(f)$. For all α we also have that the MEU-value of $f \alpha f_N$ is less than $U(f)$ on any $\Pi_c(\mathcal{K}_m)$, $m \geq N + 1$. It follows that $U(f \alpha f_N) \leq U(f)$ – contradiction.

Case (ii):

It remains to consider case (ii), where $\Pi_1^N \cap \Pi_2^N \neq \emptyset$, $\forall N$ (the logical negation of case (i) is that the intersection is non-empty for infinitely many N , but we confuse the two statements here since it is irrelevant to the argument). First, a general comment. Take one set $\Pi_c(\mathcal{K}_N)$ and let $\pi_N \in \Pi_c(\mathcal{K}_N)$ be a measure such that the value of f on $\Pi_c(\mathcal{K}_n)$ obtains on π_n . Assume for the moment that, for any π_N , we can choose a \hat{f} whose value is attained on $\Pi_c(\mathcal{K}_N)$ and at the measure π_N .²³ Abusing notation let $\pi_n \in \Pi_c(\mathcal{K}_n)$ be measures on which the MEU-value of \hat{f} is attained and take (after relabeling as needed) a (sub)sequence $\{\Pi_c(\mathcal{K}_m)\}$ along which the MEU-values of \hat{f} increase to $\sup(\text{MEU}_{\Pi_c(\mathcal{K}_n)}(\hat{f}))$. We also switch the index to m for reasons that will soon become apparent. Let $\Pi_c(\mathcal{K}_M)$ be the set on which the value \hat{f} is strictly attained. Consider two possibilities, (a) $\pi_M \in \Pi_c(\mathcal{K}_m)$ for infinitely many m or (b) $\pi_M \in \Pi_c(\mathcal{K}_m)$ for finitely many m . In case (a), find large

²³Recall that this means that there is some $\varepsilon_{\hat{f}}$ such that $\text{MEU}_{\Pi_c(\mathcal{K}_N)}(\hat{f}) > \text{MEU}_{\Pi_c(\mathcal{K}_\alpha)}(\hat{f}) + \varepsilon_{\hat{f}}$, $\forall \alpha \neq N$.

M^* such that the MEU-value of f on $\Pi_c(\mathcal{K}_{M^*})$ is strictly larger than on $\Pi_c(\mathcal{K}_M)$ and $\pi_M \in \Pi_c(\mathcal{K}_{M^*})$. It follows that, on the one hand, $E_{\pi_M} u(f) \geq \text{MEU}_{\Pi_c(\mathcal{K}_{M^*})}(f)$. On the other, $\text{MEU}_{\Pi_c(\mathcal{K}_{M^*})}(f) > \text{MEU}_{\Pi_c(\mathcal{K}_M)}(f) = E_{\pi_M} u(f)$ – contradiction, where the latter equality uses that the MEU-value of f on $\Pi_c(\mathcal{K}_N)$ obtains on π_N .

In case (b), construct a separation argument as in case (i) to obtain a contradiction. Find M^* large enough so that $\pi_M \notin \Pi_c(\mathcal{K}_{M^*})$. Choose a vector v with $\pi \cdot v > 0 > \pi_M \cdot v, \forall \pi \in \Pi_c(\mathcal{K}_{M^*})$. Put $f_\varepsilon := \hat{f} + \varepsilon \cdot v$, where we choose ε so that $\text{MEU}_{\Pi_c(\mathcal{K}_{M^*})}(f_\varepsilon) > U(\hat{f})$. Then, $U(f_\varepsilon) > U(\hat{f})$, so that (by axiom 5) $f_\varepsilon \alpha \hat{f} \succeq f, \forall \alpha$. On the other hand, we know that for all large α $\text{MEU}_{\Pi_c(\mathcal{K}_\alpha)}(f_\varepsilon \alpha \hat{f}) < U(f)$. The latter since $E_\pi u(f_\varepsilon)$ is bounded and $\text{MEU}_{\Pi_c(\mathcal{K}_\alpha)}(\hat{f}) < U(\hat{f}) - \varepsilon_{\hat{f}}$ for some $\varepsilon_{\hat{f}}$, contradiction. It follows that, in both cases, we obtain a contradiction if the value of f on $\Pi_c(\mathcal{K}_n)$ obtains on a π_n for which there is $\hat{f}, \varepsilon_{\hat{f}}$ such that $U(\hat{f})$ obtains on π_n and, moreover, $\text{MEU}_{\Pi_c(\mathcal{K}_\alpha)}(\hat{f}) < U(\hat{f}) - \varepsilon_{\hat{f}}, \forall \alpha \neq n$. We reduce to consider the case where π_n is a limit (within $\Pi_c(\mathcal{K}_n)$) of such measures. Fix the $\pi_M \in \Pi_c(\mathcal{K}_M)$ on which f attains its (local MEU) value and let $(\pi_M^l, \hat{f}_l, \varepsilon_l)$ be a sequence of triples where:

- i. \hat{f}_l attains its (global) value $U(\hat{f}_l)$ on $\Pi_c(\mathcal{K}_M)$ and at the measure π_M^l .
- ii. $U(\hat{f}_l) - \varepsilon_l > \text{MEU}_{\Pi_c(\mathcal{K}_\alpha)}(\hat{f}_l), \forall \alpha \neq M$.
- iii. $\pi_M^l \rightarrow \pi_M$.

By the preceding argument, for each (π_M^l, \hat{f}_l) there must be infinitely many $\Pi_c(\mathcal{K}_n)$ with $\pi_M^l \in \Pi_c(\mathcal{K}_n)$. Since $\text{MEU}_{\Pi_c(\mathcal{K}_M)}(f) < U(f)$, put $\delta_f := U(f) - \text{MEU}_{\Pi_c(\mathcal{K}_M)}(f)$ and find l large enough so that (by (iii))

$$|E_{\pi_M^l} u(f) - E_{\pi_M} u(f)| < \delta_f/2.$$

For this choice of l there are infinitely many $\Pi_c(\mathcal{K}_n)$ with $\pi_M^l \in \Pi_c(\mathcal{K}_n)$. Find a large enough index M^* such that $U(f) - \text{MEU}_{\Pi_c(\mathcal{K}_{M^*})}(f) < \delta_f/2$.²⁴ For this pair (l, M^*) we have, on the one hand,

$$(*) \quad E_{\pi_M^l} u(f) \geq \text{MEU}_{\Pi_c(\mathcal{K}_{M^*})}(f).$$

On the other hand,

$$\text{MEU}_{\Pi_c(\mathcal{K}_{M^*})}(f) \geq U(f) - \delta_f/2$$

and

$$U(f) - \delta_f/2 = E_{\pi_M} u(f) + \delta_f/2 > E_{\pi_M^l} u(f),$$

which together with $(*)$ yields a contradiction.

²⁴Note that such an index exists since we assumed the MEU values along the sequence of sets of measures, $\Pi_c(\mathcal{K}_n)$, increase to $U(f)$.

Remark: This proves the representation theorem under the caveat that each of the sets $\Pi_c(\mathcal{K}_\alpha)$ has the property that for any extreme point π either (i) there is an act f whose global value is its MEU value on $\Pi_c(\mathcal{K}_\alpha)$ and the value is attained on π with a cushion of size ε , or (ii) there is a sequence $(\pi^n, f_n), \pi^n \in \Pi_c(\mathcal{K}_\alpha)$ such that along this sequence the MEU-values (= global values) are attained on the measures π_n and each with a shrinking error term ε_n . We prove this property now.

Sub-Step 2b. Showing that $\Lambda^0(\beta_n)$ is finite.

Fix $\beta := \beta_n$ and, towards contradiction, say that there are infinitely many \mathcal{K}_α with $\alpha \in \Lambda^0(\beta_n)$. Let $(\Pi_c(\mathcal{K}_n), f_n)$ be the associated sequence of pairs consisting of sets of measures (supporting the cone defined from the graph \mathcal{K}_n) and acts $f_n \in \mathcal{K}_n$ whose MEU-value on $\Pi_c(\mathcal{K}_n)$ is at least β higher than the MEU-value on any other $\Pi_c(\mathcal{K}_\alpha)$. Pass wlog to a convergent subsequence and consider N_β large enough so that

- i. $|f_n - f_m| < \beta/4, \forall n, m \geq N_\beta$, and letting f_* denote the limit act,
- ii. $|f_n - f_*| \leq \beta/4, \forall n \geq N_\beta$.

This implies the MEU-value of f_n on $\Pi_c(\mathcal{K}_n)$ is within $\beta/4$ of the MEU-value of f_* and within $\beta/4$ of the MEU-value of any other f_m (for $m, n \geq N_\beta$). Therefore, the value $U(f_n)$ is within $\beta/2$ of the MEU-value of f_n on any $\Pi_c(\mathcal{K}_m)$ for any $m, n \geq N_\beta$ – contradiction.²⁵

Sub-Step 2c. Showing that $\Lambda^0(\beta_n) = \Lambda^0(\beta_m), \forall m, n$ large.

By part (a), for each act f there is an associated ε_f such that the value $U(f)$ is attained on some $\Pi_c(\mathcal{K}_\alpha)$ and, moreover, the MEU-value of f on this set of measures is bounded above the MEU-value on all other $\Pi_c(\mathcal{K}_{\alpha'})$ by at least ε . By the same (triangle inequality) argument from sub-step 2c we obtain that, for any f' such that $|f - f'| < \varepsilon/2$, the MEU-values of f, f' on *any* set $\Pi_c(\mathcal{K}_\alpha)$ can be at most $\varepsilon/2$ apart. This, in turn, implies that the values $U(f), U(f')$ are at most $\varepsilon/2$ apart. Hence, for any $f' \in B_{\varepsilon_f}(f)$ we can take $\varepsilon_{f'} = \varepsilon_f/2$. Since we have an open cover of $B(\mathcal{S})$ given by the collection, $\{B_{\varepsilon_f}(f)\}$, some finite collection $\{B_{\varepsilon_{f_1}}(f_1), \dots, B_{\varepsilon_{f_n}}(f_n)\}$ covers as well. Put $\varepsilon^* = (\min\{\varepsilon_{f_i}\}_{i=1}^n)/2$. It follows that whenever $\beta_n \leq \varepsilon^*$ we have, for all f , the MEU-value of f on the corresponding maximal set $\Pi_c(\mathcal{K}_\alpha)$ is at least ε^* above its MEU value on any other set. Hence, $\Lambda^0(\beta_n) = \Lambda^0(\beta_m), \forall m, n$ large. □

Proof of Theorem 2. By considering constant acts, we obtain that u' is an affine translate of u . We verify the uniqueness claim on the priors. This follows in three steps. First we show that each source, $\Pi_i \in \{\Pi_i\}$, contains the cone of measures

²⁵Note that this argument does not use weak uncertainty aversion. This restriction is only used in sub-step 2a.

supporting the (utility) act whose value is strictly attained on the source Π_i (we will call these ‘critical’ acts). Second, we use regularity to show that each source can be contracted (as needed) to only consist of the (closure of) union of these cones. Third, we use this characterization of each source to show that there is a unique such local MEU representation.

Step 1: Showing each Π_i contains cones of measures supporting critical acts.

Pick any $\Pi_1 \in \{\Pi_i\}$ and note that, by non-redundance, there must be an act f for which the local MEU functional associated to $(u, \{\Pi_i\})$ occurs strictly on the set Π_1 . Find ε_f small enough so that for any $g \in B_{\varepsilon_f}(f)$ with $g \succeq f$ we have (i) for the model $(u, \{\Pi_i\})$ the value at g is attained on Π_1 and (ii) for the model $(u, \{\Pi'_i\})$ the value at g is attained on one of the sets $\{\Pi'_{i_1}, \dots, \Pi'_{i_k}\}$. Consider the cone,

$$\mathcal{C}_f^{\varepsilon_f} := \{d \cdot (u(g) - u(f)) : d \geq 0, g \succeq f, g \in B_{\varepsilon_f}(f)\}.$$

Note that $\mathcal{C}_f^{\varepsilon_f}$ is closed and convex. Closure follows from continuity and the monotonicity axiom and convexity by choice of ε_f , since all acts in a small ε_f -ball of f obtain their local MEU value on the same source as f .

By the dual characterization of $\mathcal{C}_f^{\varepsilon_f}$ we find a set of (probability) measures $\Pi_f^{\varepsilon_f}$ such that

$$d \cdot (u(g) - u(f)) \in \mathcal{C}_f^{\varepsilon_f} \Leftrightarrow E_\pi u(g) \geq E_\pi u(f), \forall \pi \in \Pi_f^{\varepsilon_f}.$$

Now consider the cone,

$$\mathcal{C}_f^{\Pi_1} := \{d \cdot (u(g) - u(f)) : d \geq 0, E_\pi u(g) \geq E_\pi u(f), \forall \pi \in \Pi_1\}.$$

Since the value of the local MEU model $(u, \{\Pi_i\})$ at f is attained on the source Π_1 , we have $g \succeq f$, whenever $E_\pi u(g) \geq E_\pi u(f), \forall \pi \in \Pi_1$.

Lemma 1. $\mathcal{C}_f^{\Pi_1} \subseteq \mathcal{C}_f^{\varepsilon_f}$.

Proof. Take any $d \cdot (u(g) - u(f)) \in \mathcal{C}_f^{\Pi_1}$ and find $\hat{g} \in B_{\varepsilon_f}(f)$ (using surjectivity of the vNM index $u(\cdot)$, which we have in a neighborhood of the utility vector induced by any interior act) such that $u(\hat{g}) = u(f) + \varepsilon \cdot (u(g) - u(f))$. Note that $E_\pi(u(\hat{g}) - u(f)) = \varepsilon \cdot E_\pi(u(g) - u(f)) \geq 0, \forall \pi \in \Pi_1$. Hence, $\hat{g}_1 \succeq f, \hat{g} \in B_{\varepsilon_f}(f)$, implying that $\hat{d} \cdot (u(\hat{g}) - u(f)) \in \mathcal{C}_f^{\varepsilon_f}, \forall \hat{d} \geq 0$. Put $\hat{d} = d/\varepsilon$ and we obtain $\hat{d} \cdot (u(\hat{g}) - u(f)) = d \cdot (u(g) - u(f))$. This shows that $\mathcal{C}_f^{\Pi_1} \subseteq \mathcal{C}_f^{\varepsilon_f}, \forall \varepsilon$. \square

The particular threshold ε_f for which we are assured that the cone is closed and convex may depend on the model. However, for any $\varepsilon_1, \varepsilon_2 \leq \varepsilon_f$ with, say, $\varepsilon_1 < \varepsilon_2$ we have the containment

$$(*) \mathcal{C}_f^{\varepsilon_1} \subseteq \mathcal{C}_f^{\varepsilon_2}.$$

If $\varepsilon_1, \varepsilon_2 \leq \varepsilon_f$, then both cones are convex (and closed). Hence, taking duals we obtain $\Pi_f^{\varepsilon_1} \supseteq \Pi_f^{\varepsilon_2}$. Define $\Pi_f := \cup_{\varepsilon \leq \varepsilon_f} \Pi_f^\varepsilon$. Note that, while the particular threshold ε_f depends on the choice of model $(u, \{\Pi_i\})$ which represents \succeq , the sets Π_f *do not* depend on the representation (by $(*)$). We therefore obtain that, for any f which is critical for the source Π_1 , $\Pi_f \subseteq \Pi_1$. Doing this for every f we obtain that each collection Π_f is contained in Π_i , for any i on which the value of the model $(u, \{\Pi_i\})$ is strictly attained at the act f . Thus, for each i put

$$\widehat{\Pi}_i := \bigcup_{f \text{ critical for } \Pi_i} \Pi_f.$$

Step 2: Regularity implies $\widehat{\Pi}_i = \Pi_i$.

Sub-step 2a: We first show that if f obtains its value on the source Π_i , then the value is also attained on $\widehat{\Pi}_i$.

To show this first consider Π_1 , say, and an f which obtains its value on this source, i.e. $U(f) = \min_{\pi \in \Pi_1} E_\pi u(f)$. Find a constant act c such that $u(c) = U(f)$ and $\alpha \in (0, 1)$ such that $\alpha \cdot c + (1 - \alpha) \cdot f \in B_{\varepsilon_f}(f)$. By c -linearity and continuity, we know that $U(\alpha \cdot c + (1 - \alpha) \cdot f) = \alpha \cdot u(c) + (1 - \alpha) \cdot U(f)$. Hence, $\alpha \cdot c + (1 - \alpha) \cdot f \succeq f$. It follows that $E_\pi(\alpha \cdot u(c) + (1 - \alpha) \cdot u(f)) \geq E_\pi u(f), \forall \pi \in \Pi_f^{\varepsilon_f}$. Hence, $u(c) \geq E_\pi u(f), \forall \pi \in \Pi_f^{\varepsilon_f}$. It follows that $u(c) \geq \min_{\pi \in \widehat{\Pi}_1} E_\pi u(f)$. Since $\widehat{\Pi}_1 \subseteq \Pi_1$ and the value $U(f)$ is attained on Π_1 we obtain $\min_{\pi \in \widehat{\Pi}_1} E_\pi u(f) = u(c)$.

Sub-step 2b: Check that $u(c_f)$ is an upper bound on all $\widehat{\Pi}_i \subseteq \Pi_i$ for which $u(c_f) < \min_{\pi \in \Pi_i} E_\pi u(f)$.

This is immediate if the value $\min_{\pi \in \Pi} E_\pi u(f)$ is attained on some $\Pi_g^{\varepsilon_g} \subseteq \widehat{\Pi}_i$. Hence, consider the case where this does not happen. That is, consider the possibility that

$$u(c) < \min_{\pi \in \Pi_g^{\varepsilon_g}} E_\pi u(f), \forall \Pi_g^{\varepsilon_g} \subseteq \widehat{\Pi}_i.$$

Recall that we only include those sets $\Pi_g^{\varepsilon_g}$ in $\widehat{\Pi}_i$ for which the value $U(g)$ is strictly attained on the set Π_i . For each source i on which the value $U(f)$ is not attained, consider the set $\{\pi \in \Pi_i \setminus \widehat{\Pi}_i\}$. Refer to measures in this set as *exceptional* measures. Take a f for which $\min_{\pi \in \widehat{\Pi}_i} E_\pi u(f) > u(c_f)$ (so that f is not critical for Π_i , by the preceding sub-step). Let π_* be such an exceptional measure. We first consider the case where there is a critical act g (for $\widehat{\Pi}_i$) that also obtain its value on this exceptional measure. Put $K := \min_{\pi \in \widehat{\Pi}_i} E_\pi u(f) > u(c_f)$. Take a critical act g for Π_i which obtains its value on the exceptional measure π_* and find a pair $(\alpha, c), \alpha \in (0, 1)$ and c a constant act such that: (denote the α -mixture of g with c as $g\alpha c$)

- i. $u(c_f) < u(c) < K$,

ii. $E_\pi u(g\alpha c) < K, \forall \pi \in \widehat{\Pi}_i$.

For brevity put $\hat{g} := g\alpha c$ and find β small enough so that $f\beta\hat{g} \in B_{\varepsilon_{\hat{g}}}(\hat{g})$. Note that, since g is critical for Π_i so is \hat{g} , so that there is a $\varepsilon_{\hat{g}} > 0$ such that all elements in the ball with this radius are also critical for Π_i . By (ii), we have $E_\pi u(f\beta\hat{g}) \geq E_\pi u(\hat{g}), \forall \pi \in \Pi_{\hat{g}}^{\varepsilon_{\hat{g}}}$ – since $\Pi_{\hat{g}}^{\varepsilon_{\hat{g}}} \subseteq \widehat{\Pi}_i$. Since $f\beta\hat{g} \in B_{\varepsilon_{\hat{g}}}(\hat{g})$ and $\Pi_{\hat{g}}^{\varepsilon_{\hat{g}}}$ is dual to the cone,

$$\mathcal{C}_{\hat{g}}^{\varepsilon_{\hat{g}}} = \{d \cdot (u(h) - u(\hat{g})) : h \in B_{\varepsilon_{\hat{g}}}(\hat{g}), h \succeq \hat{g}\}$$

we obtain that, on the one hand, $f\beta\hat{g} \succeq \hat{g}$. On the other hand, we have (a) $E_{\pi_*} u(f) \leq u(c_f)$ and (b) \hat{g} obtains its value on π_* (since g does, by hypothesis). This implies that for any $\beta > 0$ we have

$$E_{\pi_*} u(f\beta\hat{g}) = \beta \cdot E_{\pi_*} u(f) + (1 - \beta) \cdot E_{\pi_*} u(\hat{g}) \leq \beta \cdot u(c_f) + (1 - \beta) \cdot u(\hat{g}) < u(\hat{g}).$$

Choosing β small enough that (a) $f\beta\hat{g} \in B_{\varepsilon_{\hat{g}}}(\hat{g})$ and (b) $f\beta\hat{g}$ is critical for Π_i , this implies that $f\beta\hat{g} \prec \hat{g}$, a contradiction.

Now consider the case where g (which is non-critical for Π_i) has the property that $E_\pi u(g) < u(c_g)$ for some exceptional measure π . By the preceding argument, we must have $E_\pi u(f) > u(c_f)$ for *all* critical acts. In this case, we amend the preceding argument as follows. For brevity, define:

1. $\min_{\pi \in \Pi_i} E_\pi u(f) = \underline{u}, E_{\pi_*} u(f) = \bar{u}, \max_{\pi \in \widehat{\Pi}_i} E_\pi u(f) = \hat{u}$.
2. $\min_{\pi \in \Pi_i} u(g) = E_{\pi_*} u(g) = \underline{v}, \min_{\pi \in \widehat{\Pi}_i} E_\pi u(g) = \bar{v}$.

By scaling g with an appropriate constant $\hat{g} := K \cdot g$ we find a constant act c such that (i) $\underline{v} < u(c) < \bar{v}$, (ii) $\bar{v} > \hat{u}$, and (iii) $u(c) > \bar{u}$. Given such c , for any $\alpha \in (0, 1)$:

$$(*) \quad \alpha \cdot \hat{u} + (1 - \alpha) \cdot u(c) < \bar{v}.$$

Replace f with the act $f\alpha c$ and note that (i) $f\alpha c$ is also critical for Π_i and (ii) for all sufficiently small β , the act $g\beta(f\alpha c)$ is in the requisite $\varepsilon_{f\alpha c}$ -neighborhood of $f\alpha c$ such that (local) uncertainty aversion holds *and* $g\beta(f\alpha c)$ is also critical for Π_i . Observe that we have $E_\pi u(g\beta(f\alpha c)) \geq E_\pi u(f\alpha c)$ if and only if $E_\pi u(g) \geq E_\pi u(f\alpha c)$. By (*) this happens for all $\pi \in \widehat{\Pi}_i$. Hence, $g\beta(f\alpha c) \succeq f\alpha c$. This holds for any pair (α, β) , where we fix any α and then choose a sufficiently small β (depending on the fixed α). We will pick a specific pair (α, β) . Choose the constant K defining \hat{g} such that there is an integer N such that²⁶

$$\text{a. } \underline{u} > \frac{2}{N} \underline{v}.$$

²⁶The constant 2 in (a), (b) below is wlog. We just need some constant $\theta > 1$. Hence, we can take (a) $\underline{u} > \frac{\theta}{N} \underline{v}$, (b) $u(c) > N \cdot \bar{u}$, replace 2 with θ in (***) and carry out the same argument verbatim. In the remarks following the proof, we show that there is always such a θ .

b. $u(c) > N \cdot \bar{u}$.

In the remarks following this argument we show that such an N always exists. Given N , pick (α, β) such that

$$(**) \frac{N}{2}\beta < \alpha < \frac{N\beta}{1 - \beta + N\beta}$$

We claim that, for (α, β) satisfying (**), we have:

$$E_{\pi_*} u(g\beta(f\alpha c)) < E_{\pi_*} u(f\alpha c).$$

The inequality can be equivalently expressed as:

$$(***) \beta \cdot \underline{v} + (1 - \beta) \cdot [\alpha \bar{u} + (1 - \alpha)u(c)] < \alpha \underline{u} + (1 - \alpha)u(c),$$

which simplifies to:

$$\beta \cdot \underline{v} + (1 - \beta)\alpha \bar{u} < \alpha \cdot \underline{u} + \beta(1 - \alpha) \cdot u(c).$$

Using $\alpha \geq \frac{N}{2}\beta$ and (a) we obtain:

$$\alpha \underline{u} \geq \frac{N}{2}\beta \underline{u} > \frac{N}{2}\beta \frac{2}{N} \underline{v} = \beta \cdot \underline{v}.$$

To show (***) it suffices then to show that:

$$(1 - \beta)\alpha \bar{u} < \beta(1 - \alpha) \cdot u(c).$$

Since $u(c) > N \cdot \bar{u}$ (by (b)) we verify that:

$$(1 - \beta)\alpha < \beta(1 - \alpha)N,$$

but this is equivalent to $\frac{\beta}{1 - \beta + N\beta} \geq \frac{\alpha}{N}$, which follows from (**).

Remarks: This concludes step 2 subject to two caveats. First, we have assumed the existence of an integer N and a scaling constant K such that (a), (b) above hold. Second, we need to check that $g\beta(f\alpha c)$ is critical for Π_i , as this is implicit in the above argument. We check the latter condition first. Since f is critical for Π_i , let ε_f be such that $\forall g \in B_{\varepsilon_f}(f)$, we have g critical for Π_i as well. For any $\alpha \in (0, 1)$ and constant act c , put $\hat{f} = f\alpha c$ and $\varepsilon_{\hat{f}} = \alpha \cdot \varepsilon_f$. Then, for any $h \in B_{\varepsilon_{\hat{f}}}(\hat{f})$ we have h is critical for Π_i . Taking the mixed act $g\beta\hat{f}$ as above, this gives the following bound between α, β that is sufficient for $g\beta\hat{f}$ to be critical for Π_i :

$$\beta|g| \leq \alpha \cdot \varepsilon_f,$$

where $|g|$ denotes the norm of the utility vector $u(g)$. This condition needs to be consistent with (**). A sufficient condition for consistency is to replace $|g|$ with $M := \max_{h \in B(\mathcal{S})} |h|$ and verify that

$$\beta \cdot M < \frac{N\beta}{1 - \beta + N\beta} \cdot \varepsilon_f.$$

This is equivalent to:

$$\frac{M}{\varepsilon_f} < \frac{N}{1 - \beta + N\beta}$$

We can (wlog) normalize so that $M \leq N$ and $N > \frac{1}{\varepsilon_f}$. The preceding inequality is equivalent to:

$$(\spadesuit) \frac{N}{M} \varepsilon_f > 1 - \beta + N\beta$$

which simplifies to: $\frac{\varepsilon_f}{M} - \beta > \frac{1-\beta}{N}$. As N, M, ε_f are fixed at the outset and we choose $N > \frac{1}{\varepsilon_f}$, we can pick β small enough so that (\spadesuit) holds. Hence, there are $\alpha, \beta \in (0, 1)$ satisfying the conditions (**), (***) and such that $g\beta(f\alpha c)$ is critical for Π_i .

It remains to check that there are (N, α, β) satisfying the conditions (**), (a), (b) (and that these can be satisfied subject to the additional condition that $M \leq N$). First, normalize utility of lotteries such that $|f| \leq 1, \forall f \in B(\mathcal{S})$. Next, given f we also have ε_f . Find N large enough so that $N > \frac{1}{\varepsilon_f}$ and replace g with the scaled (utility) act $N \cdot g$. Replace $B(\mathcal{S})$ with the space of all acts $K \cdot h, h \in B(\mathcal{S}), K \leq N$. Then, it follows that $M = N$. Call the scaled ball of acts $B_N(\mathcal{S})$. The preference obviously extends to this space. Hence, the condition that $M \leq N$ can always be arranged. We check that (**), (a), (b) can be arranged as well. For this first consider the case where $\underline{u} = \bar{u}$.²⁷ Shift g by a constant act d , i.e. replace g with $g + d$ (which is the same as first mixing $1/2 : 1/2$ with d , then scalar multiplying by 2), where we choose d to be big enough so that

$$\frac{\underline{v}}{\underline{u}} < \frac{N}{\theta}, \frac{\bar{v}}{\bar{u}} > N$$

where we have $\theta = \frac{\bar{v}}{\underline{v}} > 1$ and the dependence on $u(d)$ is via, e.g. $\bar{v} = \bar{v}_0 + u(d)$ (v_0 is the original unshifted act). Note that by mixing g with $c_{\bar{v}}$ (i.e. $u(c_{\bar{v}}) = \bar{v}$) and taking $g\alpha\bar{c}$ for α small enough we can ensure that the above displayed inequalities hold. Note that we are invoking footnote 16 here in replacing the constant 2 in (a),(b) with $\theta > 1$.

²⁷This can be arranged to be approximate equality by mixing f with some constant act \underline{c} , i.e. $f\alpha\underline{c}$, where α is sufficiently close to zero.

Step 3: Showing $(u, \Pi_i) = (u, \Pi'_i)$.

Proceed in two (sub)steps. First, we show that the set of acts which are critical, i.e. obtain the value strictly on some source Π_i for *some* representation $(u, \{\Pi_i\})$, is the same for every model $(u, \{\Pi_i\})$ which represents \succeq . In conjunction with the preceding line of argument, this immediately implies that any two sets of sources in a (regular) representation are in bijection. The second step shows that, moreover, they are identical.

Step 3a: The set of critical acts is the same for all (regular) representations $(u, \{\Pi_i\})$. We recall some notation from the proof of the representation theorem (theorem 1). We defined a symmetric relation \asymp on $B(\mathcal{S})$ and denoted by $\{\mathcal{K}\}_{\alpha \in \Lambda}$ the collection of maximal complete sub-graphs of the (un)directed graph on $B(\mathcal{S})$ induced by this relation, with a generic element \mathcal{K}_α (here we think of $\mathcal{K}_\alpha \subseteq B(\mathcal{S})$, although it is really a maximal connected component of elements of $B(\mathcal{S})$ – each joined to each other by the relation \asymp). The proof then derived a (closed, convex) set of measures, $\Pi(\mathcal{K}_\alpha)$ (we had a subscript c , for constant act c , which we suppress here for brevity). The key properties of the sets $\Pi(\mathcal{K}_\alpha)$:

1. The MEU value $U(f)$ is attained on the set $\Pi(\mathcal{K}_\alpha)$, for each $f \in \mathcal{K}_\alpha$.
2. The MEU value, $\text{MEU}_{\Pi(\mathcal{K}_\beta)}(f)$, is strictly less than $U(f)$ for each \mathcal{K}_β for which $f \notin \mathcal{K}_\beta$.
3. There are finitely many \mathcal{K}_α (for \succeq satisfying axioms 1-5).

Using these three facts, we now have the following claim. Let $\mathcal{K}_{\alpha_1}, \dots, \mathcal{K}_{\alpha_n}$ denote the finitely many maximal complete subgraphs of $B(\mathcal{S})$ (induced by \asymp).

Lemma 2. *An act f is critical for $\Pi(\mathcal{K}_{\alpha_k})$ if and only if $f \in \mathcal{K}_{\alpha_k}$ and $f \notin \mathcal{K}_{\alpha_l}, l \neq k$.*

Proof. The left-to-right direction follows from fact (3) listed above and the hypothesis that f is critical for $\Pi(\mathcal{K}_{\alpha_k})$. Consider the right-to-left direction via contraposition. Assume f is not a critical act. We know from (1) above that it obtains its MEU value $U(f)$ on each $\Pi(\mathcal{K}_{\alpha_i})$. Moreover, from fact (2) its MEU value on each $\Pi(\mathcal{K}_{\alpha_l})$ (for each l such that $f \notin \mathcal{K}_{\alpha_l}$) is strictly less than $U(f)$. It follows that there must be multiple α_i such that $f \in \mathcal{K}_{\alpha_i}$ (else, f would be a critical act). \square

Now use this to check that the set of critical acts must be the same for every (regular) model $(u, \{\Pi_i\})$ that represents \succeq . By the representation theorem (viz. its proof) there are finitely many maximal complete subgraphs, \mathcal{K}_{α_i} , say $\{\mathcal{K}_{\alpha_1}, \dots, \mathcal{K}_{\alpha_n}\}$. Note that this collection is defined just from the preference, i.e. it does not reference any specific representation $(u, \{\Pi_i\})$. By the preceding lemma, the critical acts are defined as follows:

$$\{f : \exists \alpha_i \text{ s.t. } f \in \mathcal{K}_{\alpha_i}, f \notin \mathcal{K}_{\alpha_j}, \forall j \neq i\}.$$

For any (finite) sub-collection $\{\Pi(\mathcal{K}_\alpha)\}$, call the resulting model $(u, \{\Pi(\mathcal{K}_\alpha)\})$ *basic*. The proof of theorem 1 showed that there is some finite set $\{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_m})\}$ such that for *any* basic model we have $\{\Pi(\mathcal{K}_\alpha)\} \subseteq \{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_m})\}$. We use the lemma to show there is a unique regular sub(model) of $(u, \{\Pi(\mathcal{K}_{\alpha_i})\}_{i=1}^m)$.

Lemma 3. *There is a unique regular, basic (sub)model of $\{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_n})\}$.*

Proof. The proof relies on the identification result for a (special case of) the model that we axiomatize in the paper [Chandrasekher \(2017\)](#). There we axiomatize and identify models (u, v, \mathcal{C}) , where u, v are utility functions mapping some finite X to \mathbf{R} and $\mathcal{C} = \{\mathcal{C}_i\}$ is a collection of subsets of X whose union is all of X . This triple forms a utility on menus via: $U(A) = \max_{\mathcal{C}_i} \max_{x \in (\mathcal{C}_i \cap A)_v} u(x)$. Theorem 2 in [Chandrasekher \(2017\)](#) identifies the u , the collection $\mathcal{C} = \{\mathcal{C}_i\}$, and a relation \succeq_v which underlies the v . The relation \succeq_v generates the utility $U(\cdot)$ in the sense that all completions, v' , of \succeq_v give the same utility U for the given u and \mathcal{C} . We apply this result to show uniqueness of a regular (basic) model. This procedure involves several steps, but can roughly be broken down via: (i) showing uniqueness along a sequence of finite subsets $X_n \subseteq B(\mathcal{S})$ (whose limit is dense in $B(\mathcal{S})$ and (ii) showing that the sequence of unique models (along the sequence of subsets X_n yields a unique model for the limit of the sequence. We turn to step (i).

We enlarge the space of choices, $B(\mathcal{S})$, to $B(\mathcal{S}) \times \{1, -1\}$, where we map each $f \in B(\mathcal{S})$ to $(f, 1), (f, -1)$. Take a countable dense subset of $X \subseteq B(\mathcal{S}) \times \{1, -1\}$ and let $X_n \subseteq X_{n+1}$ be an increasing sequence to X , with the additional property that whenever $f \in B(\mathcal{S})$ is such that $(f, 1) \in X_n$ (resp. $(f, -1) \in X_n$), then we also have $(f, -1) \in B(\mathcal{S})$ (resp. $(f, 1) \in X_n$). Put $X_n \equiv \{(f_1, \pm 1), \dots, (f_n, \pm 1)\}$. Define a function v_n on X_n by:

1. $v_n((f_i, 1)) > v_n(f_i, -1)$.
2. $v_n((f_i, \pm 1)) > v_n((f_j, 1)), \forall i < j$.

By choosing the values of v_n to lie in a compact subset of, say, $[0, 1]$ and choosing values of $v_n(f_n) \uparrow 1$ we obtain a consistent collection of functions $\{v_n\}$, i.e. there is a single v that recovers the v_n by restricting v to X_n . Define a collection of (a collection of) sets $\mathcal{C}(n) \equiv \{\mathcal{C}_i(n)\}_{i=1}^N$ as follows. Let the integer N denote the size of the collection $\{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_N})\}$. For each $\Pi(\mathcal{K}_{\alpha_i})$ define a set $\mathcal{C}_i(n)$ as follows:

$$\begin{aligned} (f_j, 1) \text{ (resp. } (f_j, -1)) &\in \mathcal{C}_i(n) \\ &\Leftrightarrow \text{MEU}_{\Pi(\mathcal{K}_{\alpha_i})}(f_j) = U(f_j) \text{ (resp. } \text{MEU}_{\Pi(\mathcal{K}_{\alpha_i})}(f_j) < U(f_j)). \end{aligned}$$

Take $u_n = -v_n$ and consider the model $(u_n, v_n, \mathcal{C}(n))$. There is a unique sub-model of the collection \mathcal{C}_n that is identified within the class of models (u, v, \mathcal{C}) (on a prize domain of X_n) which represent the same preference on menus as the utility $U(\cdot)$

defined by the model $(u_n, v_n, \mathcal{C}(n))$. Let $\mathcal{C}'(n)$ denote the identified sub-collection. Let \mathcal{M}_n denote the full space of (non-empty) menus on the prize set X_n and let $X_n^+ \subseteq X_n$ denote the subspace of menus of the form $\{(f_j, 1), (f_j, -1)\}$. Identification of $\mathcal{C}'(n)$ implies that, for each set $\mathcal{C}_i(n)$, we have some menu $M(i)$ such that the value $U(M(i))$ obtains (strictly) on the set $\mathcal{C}_i(n)$. By definition of the function v_n and the sets $\mathcal{C}_i(n)$, the menu $M(i)$ can (wlog) be taken to be an element of \mathcal{M}_n^+ . Hence, each set $\mathcal{C}_i(n) \in \mathcal{C}'(n)$ is (within the collection $\mathcal{C}'(n)$) critical for some act $f_j \in \{f_1, \dots, f_n\}$. Since the functions v_n, u_n extend to a single v, u we obtain a sequence of models $(u, v, \mathcal{C}'(n))$ with the property that, for each n , every $\mathcal{C}_i(n) \in \mathcal{C}'(n)$ contains some act f_j for which the set $\mathcal{C}_i(n)$ (i.e. $\Pi(\mathcal{K}_{\alpha_i})$) is critical.

Now we turn to step (ii). For each collection $\mathcal{C}'(n)$ we associate a collection of indicator variables $\{\mathbf{1}'_i(n)\}_{i=1}^N$, where we put $\mathbf{1}'_i(n) = 1$ if and only if $\mathcal{C}_i(n) \in \mathcal{C}'(n)$. Since there only finitely many sub-collection of indicators, there is a subsequence of n , say n_i , along which the sequence of indicators becomes constant. Since the sets $X_n E$ are increasing, this implies that every convergent subsequence of indicators has the same limit set of indicators. To see this, take any two convergent subsequences of indicators $\{\mathbf{1}'_i(n_k)\}, \{\mathbf{1}'_i(m_j)\}$. Pass to a large enough M such that both sequences constant. Pick any $n_k, m_j \geq M$ and say that $n_k \leq m_j$. Since the sets X_n are increasing, the collection $\{\mathcal{C}'_i(m_j)\}$ restricts to give a representation on $2^{X_{n_k}} \setminus \emptyset$. Passing to a sharp (this is the analog of regularity in the discrete choice context, see [Chandrasekher \(2017\)](#) for the definition) sub-model and applying uniqueness we obtain an inclusion on the collection of indicators, $\{\mathbf{1}'_i(n_k)\} \subseteq \{\mathbf{1}'_i(m_j)\}$. Reverse the argument, choosing some $n_k \geq m_j \geq M$ and find that $\{\mathbf{1}'_i(m_j)\} \subseteq \{\mathbf{1}'_i(n_k)\}$. Since both collections are constant, it follows that the limit set of indicators is well-defined. Let $\{\mathbf{1}^*_i\}_{i=1}^N$ denote the limit set of indicators and consider the associated sets, $\{\Pi(\mathcal{K}_{\alpha_{i_1}}), \dots, \Pi(\mathcal{K}_{\alpha_{i_K}})\} (\subseteq \{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_N})\})$. Since X is dense in $B(\mathcal{S}) \times \{1, -1\}$, the local MEU model using this sets gives a representation of \succeq .

We use this to a unique basic sub-model of the set $\{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_N})\}$. Fixing a dense set X constructs one (unique) set which is recovered as the limit of the (constant) sequence of indicators. Consider any alternative basic model, $\{\Pi(\mathcal{K}_{\alpha_{j_1}}), \Pi(\mathcal{K}_{\alpha_{j_2}}), \dots, \Pi(\mathcal{K}_{\alpha_{j_L}})\}$. Extract from this a sequence of sub-collections $\widehat{\mathcal{C}}(n)$, where we consider the increasing sequence X_n as defined above and put

$$f_j \in X_n \in \widehat{\mathcal{C}}_i(n) \text{ if and only if } U(f_j) = \text{MEU}_{\Pi(\mathcal{K}_{\alpha_i})}(f_j).$$

Since the collection $\{\Pi(\mathcal{K}_{\alpha_{j_1}}), \dots, \Pi(\mathcal{K}_{\alpha_{j_L}})\}$ is assumed to be basic, the collection of indicators (associated to the sequence $\widehat{\mathcal{C}}(n)$) must be eventually constant and equal to $\{\mathbf{1}_{\alpha_{j_1}}, \dots, \mathbf{1}_{\alpha_{j_L}}\}$. On the other hand, by the argument in the preceding paragraph we must have an equality of indicators

$$\{\mathbf{1}_{\alpha_{i_1}}, \mathbf{1}_{\alpha_{i_2}}, \dots, \mathbf{1}_{\alpha_{i_K}}\} = \{\mathbf{1}_{\alpha_{j_1}}, \dots, \mathbf{1}_{\alpha_{j_L}}\}.$$

It follows that the two (basic) collections of sets:

1. $\{\Pi(\mathcal{K}_{\alpha_{i_1}}), \dots, \Pi(\mathcal{K}_{\alpha_{i_K}})\}$, and
2. $\{\Pi(\mathcal{K}_{\alpha_{j_1}}), \Pi(\mathcal{K}_{\alpha_{j_2}}), \dots, \Pi(\mathcal{K}_{\alpha_{j_L}})\}$,

are the same. This completes the argument that there is a unique basic sub-model of the full collection $\{\Pi(\mathcal{K}_{\alpha_1}), \dots, \Pi(\mathcal{K}_{\alpha_N})\}$. \square

We now complete the argument that the set of critical acts is the same for every representation $(u, \{\Pi_i\})$ representing \succeq . Fix the unique basic sub-model, call it $(u, \{\Pi_i^*\})$ found in the lemma. Note that, while this model was recovered upon considering a limit of unique models along a pre-selected dense sequence X_n of acts, the collection Π_i^* is independent of this limit. Take any other regular model $(u, \{\Pi'_i\})$. Note that we can, for any pre-selected dense (increasing) collection X_n , apply the argument of the preceding lemma to the grand collection $\{\{\Pi_i^*\}, \{\Pi'_i\}\}$. Take any f critical for the model $(u, \{\Pi'_i\})$ and consider a dense increasing sequence X_n where we introduce f at some X_n . Note that we have two sharp sequences of sub-collections coming out of these two models, call them $\mathcal{C}^*(n), \mathcal{C}'(n)$. By uniqueness, for all large N these two sets must agree – implying that, since f is in a unique $\mathcal{C}'_i(n)$, f is in a unique $\mathcal{C}^*_i(n)$ – implying that f is critical for $(u, \{\Pi_i^*\})$.

Step 3b: The set of critical acts uniquely determines the (regular) model.

Towards contradiction, allege two distinct (regular) models $(u, \{\Pi_i\}), (u, \{\Pi'_i\})$. By step 1, each set Π_i (resp. Π'_i) is the (closure of) the union of the sets, Π_f , for acts f which are critical on the set Π_i (resp. Π'_i). Hence, there must be a pair of critical acts (f, g) which are critical for a common $\Pi_1 \in \{\Pi_i\}$, yet which are critical for different sets $\Pi'_{i_1}, \Pi'_{i_2} \in \{\Pi'_i\}$ (say f is critical for Π'_{i_1} and g for Π'_{i_2}). Let $U(\cdot)$ denote the local MEU utility induced by the pair $(u, \{\Pi'_{i_1}\})$ and let $\mathcal{I} := \{i_1, i_2, \dots, i_n\}$ be a list of the sets Π'_{i_j} comprising the model $(u, \{\Pi'_{i_j}\})$. Consider the map $\phi : [0, 1] \rightarrow 2^{\mathcal{I}}$ defined by, $\alpha \mapsto \{j \in \mathcal{I} : U(f\alpha g) = \text{MEU}_{\Pi'_{i_j}}(f)\}$. We claim that there must be some pair $i, j \in \mathcal{I}$ and some α for which $i, j \in \phi(\alpha)$. To see this, note that for small α , i.e. nearly all weight on g , the function is negative and for large α it is positive. Hence, the function $\phi(\cdot)$ is not constant. If, towards contradiction, there is no α for which $\phi(\alpha)$ contains two or more indices, then – by continuity of the MEU functionals – the inverse images $\phi^{-1}(i)$ are either empty or disjoint, closed intervals in $[0, 1]$ whose union *partitions* $[0, 1]$. But now find two indices i_*, j_* such that $\sup(\phi^{-1}(i_*)) = \inf(\phi^{-1}(j_*))$ and we obtain, by continuity of $U(\cdot)$, $U(\sup(\phi^{-1}(i_*))) = U(\inf(\phi^{-1}(j_*)))$. On the other hand, setting the common value as α_* , we must have $\phi(\alpha_*) = i_* = j_*$ – contradiction. Hence, there is some α_* for which the associated mixture act $f\alpha_*g$ is not a critical act for the model $(u, \{\Pi'_i\})$. We now check that for any α , the mixture $f\alpha g$

is critical for the model $(u, \{\Pi_i\})$, giving a contradiction to step 3a.

Consider the model $(u, \{\Pi_i\})$, where f, g are critical for a common source, say Π_1 . Let $\{\pi_i^f\}_{i \neq 1}, \{\pi_i^g\}_{i \neq 1}$ denote the (resp.) measures in source Π_i which attain the MEU-value of f (resp. g). For the model $(u, \{\Pi'_i\})$ we put Π'_{i_1}, Π'_{i_2} as the (resp.) sources on which the values $U(f), U(g)$ are attained and let $\{\pi_{i_j}^f\}_{j \neq 1}, \{\pi_{i_j}^g\}_{j \neq 2}$ be (resp.) measures in Π'_{i_j} on which the MEU-values of f, g are attained. Consider, for both f, g , the sets $\mathcal{K}_1^f := \Pi_1 \cup \Pi'_{i_1}, \mathcal{K}_2^f := \text{co}(\{\pi_i^f\}_{i \neq 1}, \{\pi_{i_j}^f\}_{j \neq 1})$ (and similarly define $\mathcal{K}_1^g, \mathcal{K}_2^g$). Since the models $(u, \{\Pi_i\}), (u, \{\Pi'_i\})$ represent the same menu preference, they must induce the same utility function on acts (as the u 's are the same). This implies that

$$\mathcal{K}_1^f \cap \mathcal{K}_2^f = \emptyset,$$

and similarly $\mathcal{K}_1^g \cap \mathcal{K}_2^g = \emptyset$. Hence, we can choose vectors v_f, v_g such that:

$$(\clubsuit) \quad \pi \cdot v_f > 0 > \pi' \cdot v_f, \forall \pi \in \Pi_1, \forall \pi' \in \mathcal{K}_2^f$$

(and similarly for g). For any positive scalars t, t' we therefore have,

- i. $E_\pi(u(f) + tv_f) \geq E_{\pi'}(u(f) + tv_f), \forall \pi \in \Pi_1, \forall \pi' \in \mathcal{K}_2^f$.
- ii. $E_\pi(u(g) + t'v_g) \geq E_{\pi'}(u(g) + t'v_g), \forall \pi \in \Pi_1, \forall \pi' \in \mathcal{K}_2^g$.

For any positive t, t' we let $f_t, g_{t'}$ be the associated perturbed (and scaled) whose utility values give $u(f) + tv_f, u(g) + t'v_g$. Note that the inequalities imply that the acts $f_t, g_{t'}$ are (resp.) critical for Π'_{i_1}, Π'_{i_2} . We now choose t, t' such that, for any α , the act $f_t \alpha g_{t'}$ is critical for Π_1 . For each source $i \neq 1$ we check that there is a fixed (t, t') chosen at the outset such that

$$\min(E_{\pi_i^f} u(f_t \alpha g_{t'}), E_{\pi_i^g} u(f_t \alpha g_{t'})) \stackrel{\text{I}}{<} \text{MEU}_{\Pi_1}(f \alpha g) \stackrel{\text{II}}{\leq} \text{MEU}_{\Pi_1}(f_t \alpha g_{t'}).$$

The latter inequality (II) follows trivially from (\clubsuit) above. We check the former inequality (I) holds.

We show that, for an appropriate choice of (t, t') (and after scaling v_f, v_g as needed), the functions

$$\Phi_i(\alpha) := \min\{E_{\pi_i^f}(\alpha \cdot tv_f + (1 - \alpha) \cdot t'v_g), E_{\pi_i^g}(\alpha \cdot tv_f + (1 - \alpha) \cdot t'v_g)\}$$

are negative for all i . For brevity, put $\Phi_i^f(\alpha) = E_{\pi_i^f}(\alpha \cdot tv_f + (1 - \alpha) \cdot t'v_g), \Phi_i^g(\alpha) = E_{\pi_i^g}(\alpha \cdot tv_f + (1 - \alpha) \cdot t'v_g)$ and let $|E_{\pi_i^f} tv_f| =: E_{\pi_i^f}^+ tv_f, |E_{\pi_i^g} t'v_g| =: E_{\pi_i^g}^+ t'v_g$. Note that:

- a. $\Phi_i^f(\alpha) \geq 0 \Leftrightarrow t/t' \geq (1 - \alpha) E_{\pi_i^f}^+ v_g / \alpha E_{\pi_i^f}^+ v_f$.

$$\text{b. } \Phi_i^g(\alpha) \geq 0 \Leftrightarrow t/t' \leq (1 - \alpha)E_{\pi_i^g}^+v_g/\alpha E_{\pi_i^g}v_f.$$

Note that if there is an α_* such that $\Phi_i^f(\alpha_*) \geq 0$ and $\Phi_i^g(\alpha_*) \geq 0$, then – since this implies $\Phi_i^f(\alpha) \geq 0, \forall \alpha \geq \alpha_*$ and $\Phi_i^g(\alpha) \geq 0, \forall \alpha \leq \alpha_*$ – we have $\Phi_i(\alpha) \leq 0, \forall \alpha$. Hence, for each i , consider the (possibly empty) interval

$$[E_{\pi_i^f}v_g/E_{\pi_i^g}^+v_g, E_{\pi_i^f}^+v_f/E_{\pi_i^g}v_f].$$

If this is a non-empty interval for each i , then we may pick any t, t' (positive) and obtain that $\Phi_i(\alpha) \leq 0, \forall i$. If this interval is empty for some i proceed as follows. By choice of (v_f, v_g) (implicitly, the entire class of pairs (cv_f, dv_g)) we have $\Phi_i^f(\alpha) > 0, \forall \alpha$ close to unity and $\Phi_i^g(\beta) > 0, \forall \beta$ close to 0. Find $\bar{\alpha}, \underline{\beta}$ such that for all i , the intervals

$$(**)_i : \left[\frac{(1 - \bar{\alpha}) \cdot E_{\pi_i^g}v_f}{\bar{\alpha} \cdot E_{\pi_i^f}^+v_f}, \frac{(1 - \underline{\beta}) \cdot E_{\pi_i^g}^+v_g}{\underline{\beta} \cdot E_{\pi_i^f}v_g} \right]$$

are non-empty. Pick any t, t' such that the ratio t/t' falls in these intervals for all i . Now scale v_f, v_g to (resp.) (cv_f, dv_g) such that

$$E_{\pi_i^f}[\alpha \cdot cv_f + (1 - \alpha) \cdot dv_g] < 0, \forall i, \forall \alpha \in (\underline{\beta}, \bar{\alpha}).$$

Note that the intervals of admissible ratios t/t' defined by (i),(ii) above are constant on pairs (cv_f, dv_g) . In sum, choosing t, t' so that the ratio t/t' lies in the interval $(**)_i$ for all i , and scaling (v_f, v_g) to (cv_f, dv_g) (and replacing the original pair with the scaled up v_f, v_g) we put f_t equal to an act with induced utility $u(f_t) = u(f) + t \cdot v_f$ and, similarly, $g_{t'}$ such that $u(g_{t'}) = u(g) + t' \cdot v_g$. It follows that, for any mixture $f_t \alpha g_{t'}$, the value of this act is strictly obtained on the source Π_1 . \square

References

- ABDELLAOUI, M., A. BAILLON, L. PLACIDO, AND P. P. WAKKER (2011): “The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation,” *American Economic Review*, 101, 695–723.
- CHANDRASEKHER, M. (2017): “Informal commitment in planner-doer games,” *Mimeo*.
- CHEW, S. H. AND J. S. SAGI (2008): “Small worlds: Modeling attitudes towards sources of uncertainty,” *Journal of Economic Theory*, 139, 1–24.
- ELLSBERG, D. (1961): “Risk, Ambiguity, and the Savage Axioms,” *The Quarterly Journal of Economics*, 75, 643–669.
- ERGIN, H. I. AND F. GUL (2009): “A Theory of Subjective Compound Lotteries,” *Journal of Economic Theory*, 144, 899–929.
- FOX, C. AND A. TVERSKY (1995): “Weighing Risk and Uncertainty,” *Psychological Review*, 102, 269–283.
- GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): “Differentiating Ambiguity and Ambiguity Attitude,” *Journal of Economic Theory*, 118, 133–173.
- GILBOA, I. AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- GUL, F. AND W. PESENDORFER (2015): “Hurwitz expected utility and subjective sources,” *Journal of Economic Theory*, 159, 465–488.
- HEATH, F. AND A. TVERSKY (1991): “Preference and belief: Ambiguity and competence in choice under uncertainty,” *Journal of Risk and Uncertainty*, 4, 4–28.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): “A Smooth Model of Decision Making Under Uncertainty,” *Econometrica*, 73, 1849–1892.
- MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, 74, 1447–1498.
- SCHMEIDLER, D. (1989): “Subjective Probability and Expected Utility without Additivity,” *Econometrica*, 57, 571–587.