Recursive Expected Uncertain Utility

Simon Grant

Abstract

I extend Gul and Pesendorfer's expected uncertain utility (EUU) to decision problems in which a signal is observed before choosing an act from a menu. Outside subjective expected utility, tree-independence, within-event consequentialism, and dynamic consistency cannot generally coexist. I preserve the EUU representation of risk and ambiguity and recover within-event consequentialism together with dynamic consistency through a unique *uniform conditional belief system* for every signal, including signals with non-measurable cells. For measurable signals, the model reduces to static EUU; otherwise, ex ante evaluation depends on how information is revealed. A neo-additive subclass yields parsimonious "neo-EU" forms that accommodate Machina-style reflection.

Keywords: ambiguity, dynamic choice, conditional beliefs, dynamic consistency, consequentialism, neo-additive acts

JEL: D80, D81

1. Introduction

Gul and Pesendorfer (2014) provide behavioral foundations for a (static) model of choice under uncertainty that they dub Expected Uncertain Utility (EUU). When choosing between Savage acts, an EUU maximizer is endowed with a prior μ defined on a σ-algebra $\mathcal{E}$ of *ideal events*. Gul and Pesendorfer interpret ideal events as those aspects of uncertainty that the decision maker can quantify precisely with a probability.

To evaluate an arbitrary act f, he first forms *ideal* (that is, μ-measurable) greatest lower and least upper bounds $\underline{f}^{\mu}$ and $\overline{f}^{\mu}$ that bound the range of possible outcomes he cannot quantify precisely with his prior.[1] The *expected*

[1] That is, $\underline{f}^{\mu}$ (respectively, $\overline{f}^{\mu}$) is the largest (respectively, smallest) measurable function

uncertain utility of f is then given by

$$\int u\left(\underline{f}^{\mu}, \overline{f}^{\mu}\right) d\mu\,, \tag{1}$$

where the *interval utility* $u(x, y)$ is the utility assigned to an unquantifiable uncertain prospect with prizes lying between x and y.

Gul and Pesendorfer show their model can account for many observed deviations from subjective expected utility documented in Ellsberg-style experiments, including those considered by Machina (2009).

A Dynamic Setting: Information Decision Problems

Since applications of decision theoretic models in economics frequently involve the sequential resolution of uncertainty, this raises the question: how should a static model like EUU be extended to a dynamic setting? The approach I take is to consider an EUU maximizer facing what I refer to as an *information decision problem* (IDP). As in Gul and Pesendorfer's (2014) framework, the objects of choice are still (Savage) acts. In addition to specifying a menu of available acts, each IDP includes a signal (structure) σ, modeled as a function from the state space to a space of (conceivable) signal realizations.[2] As the EUU maximizer's actual choice from the menu is made *after* he observes the realization of the signal, *ex ante* he formulates a *plan (of action)* that assigns to each possible realization of the signal an act from the menu. Each feasible plan may be associated with a *signal-act*, a pair comprising the IDP's signal and the act induced by the plan. The EUU maximizer's (ex ante) optimal feasible plan can then be determined from his preference relation $\succsim$ over the set of signal-acts induced by feasible plans. An IDP with n possible signal realizations and m available acts has m^n feasible plans. If each plan induces a distinct act, determining an optimal plan directly from $\succsim$ requires $m^n - 1$ preference comparisons.

The following example makes the distinction between feasible plans and induced (reduced-form) acts explicit, and also illustrates the combinatorial growth in the number of feasible plans as the signal is refined.

that is smaller (respectively, larger) than f.

[2] I assume each menu is finite as is the range of each signal.

Example 1. *Let* $\Omega = [0,1]$ *and* $X = \{0,1\}$*. Let the signal* σ *have two realizations* s', s'' *with information cells*

$$\sigma^{-1}(s') = [0, \tfrac{1}{2}), \qquad \sigma^{-1}(s'') = [\tfrac{1}{2}, 1].$$

Consider the menu $M = \{f, g, h\}$ *where*

$$f(\omega) = \mathbf{1}\{\omega < \tfrac{1}{2}\}, \qquad g(\omega) \equiv 0, \qquad h(\omega) = \mathbf{1}\{\omega \geqslant \tfrac{1}{2}\}.$$

A plan *assigns to each realization of* σ *an act from* M*. Hence there are* $3^2 = 9$ *feasible plans. Each plan induces a reduced-form act obtained by applying the prescribed act on the realized information cell. For instance, the plan "choose* f *if* s'*, else choose* g*" induces the act*

$$\omega \mapsto \begin{cases} f(\omega), & \omega \in [0, \tfrac{1}{2}), \\ g(\omega), & \omega \in [\tfrac{1}{2}, 1], \end{cases}$$

which coincides with f *(since* $f(\omega) = 0$ *on* $[\tfrac{1}{2}, 1]$*). Likewise, the plan "always choose* f*" also induces* f*. Thus distinct plans can induce the same reduced-form act, even though the agent chooses contingently on the signal. By contrast, the plan "choose* f *if* s'*, else choose* h*" induces the constant act* $\mathbf{1}$*.*

More generally, if a signal σ has n realizations and there are m alternatives in a menu M, then there are m^n feasible plans in the IDP (σ, M). This combinatorial growth motivates using signal-realization contingent (*interim*) preferences to compute optimal plans recursively rather than comparing all feasible plans directly.

Three desiderata, and where dynamic non-EU models must "pay"

The goal of this paper is to provide a dynamic extension of EUU that preserves three desiderata that are natural in a sequential setting. First, the interim evaluation used after observing a signal realization should be *model consistent* with EUU: conditional beliefs should be represented by a (suitably interpreted) conditional version of the prior, and risk and uncertainty preferences should continue to be encoded by the same interval utility u. Second, interim preferences should be *consequentialist*: conditional rankings given a realized information cell should depend only on payoffs on that cell (and not on what would have been chosen in counterfactual information cells).[3] Third, the induced contingent plans should be *dynamically consistent*: an ex ante optimal plan should not generate a strict incentive to deviate after any realized signal.

Because terms like *recursive*, *consequentialist*, and *dynamically consistent* are used differently across the literature, it is helpful to separate these desiderata from a further, distinct EU-benchmark property: *tree-independence*. To that end, I distinguish three properties relative to expected utility with Bayesian updating:

1. **Tree-independence (ex ante).** Ex ante rankings depend only on the induced (reduced-form) act, not on the signal structure used to describe the resolution of uncertainty.
2. **Within-event consequentialism (interim).** Conditional rankings at a realized information cell depend only on the restriction of acts to that cell, not on payoffs outside it.
3. **Dynamic consistency.** Any ex ante optimal contingent plan remains optimal to carry out at every reached information cell.

Once one departs from subjective expected utility, these three properties cannot in general all be retained. REUU preserves within-event consequentialism and dynamic consistency, but it pays by sacrificing tree-independence: when the signal is not adapted to ideal events, ex ante evaluation depends on the signal structure. Equivalently, outside the ideal-signal case, REUU does not preserve signal-independence of ex ante rankings. For ideal signals, this tension disappears: tree-independence is restored, and the dynamic model collapses to static EUU.

Ideal Signals

Suppose the signal σ is ideal. Then for each possible realization s, the event $\sigma^{-1}(s)$ is in $\mathcal{E}$; that is, the associated information cell is ideal. Since the EUU maximizer does not perceive there to be any ambiguity with regard to the resolution of uncertainty of an ideal signal, it is natural to require that the restriction of $\succsim$ to signal-acts involving ideal signals *agrees* with his static

[3] Terminology varies: some authors use "consequentialism" to refer to a property of the decision tree or to a stronger invariance with respect to the description of the information structure. Throughout, I use "consequentialist" in the within-event sense spelled out here.

preferences. Thus, applying the representation of static preferences from (1), for any pair of ideal signals σ and $\widehat{\sigma}$ and any pair of acts f and $\widehat{f}$,

$$(\sigma, f) \succsim \left(\widehat{\sigma}, \widehat{f}\right) \iff \int u\left(\underline{f}^{\mu}, \overline{f}^{\mu}\right) d\mu \geqslant \int u\left(\underline{\widehat{f}}^{\mu}, \overline{\widehat{f}}^{\mu}\right) d\mu\,. \tag{2}$$

Hence, the *indirect* utility the EUU maximizer assigns to an IDP with an ideal signal—with which he can compare the desirability of this IDP with other IDPs with ideal signals—is simply the static expected uncertain utility of the act induced by an optimal plan.

In addition, an optimal plan can be constructed *recursively.* For each possible realization s of σ, let $\succsim_s^\sigma$ denote an (interim) preference relation over acts derived from $\succsim$ by setting

$$f \succsim_s^\sigma \widehat{f} \text{ if } (\sigma, f) \succsim (\sigma, \widehat{f}_{\sigma^{-1}(s)} f)\,, \tag{3}$$

where $\widehat{f}_{\sigma^{-1}(s)} f$ denotes the act that agrees with $\widehat{f}$ on the event $\sigma^{-1}(s)$ and with f elsewhere. Straightforward algebraic manipulation of (1) yields the representation

$$V_s^\sigma(f) = \int_{\sigma^{-1}(s)} u\left(\underline{f}^{\mu_s^\sigma}, \overline{f}^{\mu_s^\sigma}\right) d\mu_s^\sigma\,, \tag{4}$$

where μ_s^σ is the posterior belief over the (relative) sigma-algebra $\mathcal{E}_s^\sigma := \{E \cap \sigma^{-1}(s) \colon E \in \mathcal{E}\}$, derived from μ by setting

$$\mu_s^\sigma\big(E \cap \sigma^{-1}(s)\big) := \frac{\mu(E \cap \sigma^{-1}(s))}{\mu(\sigma^{-1}(s))}\,,$$

and where $\underline{f}^{\mu_s^\sigma}$ (respectively, $\overline{f}^{\mu_s^\sigma}$) is the largest (respectively, smallest) function measurable with respect to μ_s^σ that is smaller (respectively, larger) than f on the information cell $\sigma^{-1}(s)$ and agrees with f outside $\sigma^{-1}(s)$.[4] Given these signal-realization contingent (interim) preferences, a plan is optimal if and only if for each realization the act assigned by the plan for that realization is, according to these interim preferences, undominated by every other available act in the menu M. Hence if the signal has n possible realizations

[4] Notice that, since $\sigma^{-1}(s) \in \mathcal{E}$, $\underline{f}^{\mu_s^\sigma}$ (respectively, $\overline{f}^{\mu_s^\sigma}$) agrees with $\underline{f}^{\mu}$ (respectively, $\overline{f}^{\mu}$) on the information cell $\sigma^{-1}(s)$. However, this will turn out not to be the case in the sequel when I extend the analysis to signals with information cells that need not be ideal.

and the menu has m available acts then finding an optimal plan requires only (at most) $n \times (m-1)$ (interim) preference comparisons.[5]

Summary (ideal signals). In the benchmark case of ideal signals, the induced interim preferences $\succsim_s^\sigma$ are EUU preferences with Bayesian posteriors and the same interval utility u. Consequently, they satisfy the desiderata of model consistency, within-event consequentialism, and dynamic consistency. Moreover, because no non-ideal residual arises when every information cell $\sigma^{-1}(s)$ is ideal, ex ante evaluation in this benchmark is tree-independent and coincides with the static EUU functional.

More explicitly, if we let v denote the Bernoulli utility derived from the interval utility by setting $v(x) := u(x, x)$, then it follows from (4) that $f \sim_s^\sigma v^{-1} \circ V_s^\sigma(f)$. That is, the outcome $v^{-1} \circ V_s^\sigma(f)$ is the *conditional certainty equivalent* of the act f given s is the realization of the signal σ. Now define $c_s^\sigma(f) := v^{-1}(V_s^\sigma(f))$ and construct an act f_σ adapted to σ by setting its restriction on each information cell $\sigma^{-1}(s)$ equal to $c_s^\sigma(f)$. In the ideal-signal benchmark, this yields the following convenient decomposition of the ex ante value:

$$V(\sigma, f) = \int u\left(\underline{f}^\mu, \overline{f}^\mu\right) d\mu = \sum_{s:\, \mu(\sigma^{-1}(s))>0} V_s^\sigma(f)\, \mu(\sigma^{-1}(s))$$

$$= \sum_{s:\, \mu(\sigma^{-1}(s))>0} v(c_s^\sigma(f))\, \mu(\sigma^{-1}(s)) = \int v \circ f_\sigma \, d\mu\,.$$

Extension to Arbitrary Signals

The model developed in section 2 is one in which this recursive evaluation can be adapted to work for *all* IDPs. In particular, I retain the same interim preference relation $\succsim_s^\sigma$ defined in (3) and the same certainty-equivalent construction $c_s^\sigma(f) = v^{-1}(V_s^\sigma(f))$ used above to build the σ-adapted act f_σ. The novelty is that when $\sigma^{-1}(s) \notin \mathcal{E}$ the posterior belief μ_s^σ is not obtained by the standard Bayesian update, but instead is generated by a modified updating rule that extends Bayesian conditioning to non-ideal information cells. Formally, for each signal σ and each of its possible realizations s, the EUU maximizer attaches a posterior belief μ_s^σ over the sigma-algebra

[5] The number of comparisons could be fewer should one or more of the acts agree with each other on one or more of the signal's information cells.

$\mathcal{E}^\sigma_s = \{E \cap \sigma^{-1}(s) \colon E \in \mathcal{E}\}$ that, in conjunction with the interval utility u from (1), characterizes an EUU representation of the form (4) for the interim preferences $\succsim^\sigma_s$.

Using these signal-realization contingent functionals the EUU maximizer can compute an optimal plan by assigning to each realization an act from the menu that maximizes the corresponding conditional expected uncertain utility. Apart from the modification to the updating rule needed to handle non-ideal information cells, these interim preferences continue to satisfy model consistency, (within-event) consequentialism, and dynamic consistency. Furthermore, I show that the utility the EUU maximizer assigns an arbitrary signal-act (σ, f) may be expressed as the expected *uncertain* utility of an act f_σ adapted to σ and defined by setting its restriction on each information cell $\sigma^{-1}(s)$ equal to $v^{-1} \circ V^\sigma_s(f)$. That is,

$$V(\sigma, f) = \int u\Big(\underline{f_\sigma}^\mu, \overline{f}^\mu_\sigma\Big)\; d\mu\,. \tag{5}$$

Correspondingly, if we consider an arbitrary IDP with signal σ and menu M for which the feasible plan $\phi \colon \{s : \mu(\sigma^{-1}(s)) > 0\} \to M$ is optimal, then the *indirect* utility assigned by the EUU maximizer to this IDP, with which he can compare its desirability with respect to *any* other IDP, may be expressed as the expected *uncertain* utility of an act ϕ_σ adapted to σ that is defined by setting its restriction on each information cell $\sigma^{-1}(s)$ equal to $v^{-1} \circ V^\sigma_s(\phi(s))$.

As I noted above, for ideal signals, ex ante evaluation collapses to static EUU (cf. (2)). Outside this benchmark class, however, the analogous equality need not hold.[6] I do not view this as a defect, but rather as the natural implication of how ambiguity is perceived in a sequential environment. Channelling the insights of Frisch and Baron (1988), ambiguity about a realized information cell arises from missing information that is relevant and could in principle be known. Once the information cell is known to have obtained, the only remaining ambiguity concerns the probabilities of subevents of that cell; ambiguity that would have been faced under other, now counterfactual, cells is no longer decision relevant *interim*. However, that borne (counterfactual) ambiguity can matter *ex ante*, before the realization is learned, and

[6] Indeed, the equality only holds for a signal-act involving a signal that is not adapted to $\mathcal{E}$ when the act is constant on every information cell associated with the signal that is not ideal.

its impact is precisely what is captured by the dependence of $V(\sigma, f)$ on the signal structure in (5).

Roadmap

Section 2 formalizes information decision problems and derives the REUU evaluation of signal-acts, first for ideal signals and then for arbitrary signals via the uniform conditional belief system. Section 3 provides the axiomatic foundations linking ex ante and interim evaluations and yielding the representation theorem. Section 4 studies the neo-additive subclass, giving a parsimonious "neo-EU" representation. Section 5 situates the approach relative to existing decision-theoretic and statistical updating literatures.

2. The Model

Adopting Gul and Pesendorfer's setting of purely subjective uncertainty, Ω (with generic element ω) denotes the state space, subsets of which are referred to as events. The non-degenerate interval $X = [\ell\,, m]$ comprises the set of (monetary) outcomes and $\mathcal{F}$, the set of mappings from Ω to X, is the set of (Savage) acts, the objects of choice. Each outcome $x \in X$ is also identified with the (constant) act f in which $f(\omega) = x$ for all ω. And with further (albeit fairly standard) abuse of notation, X will also refer to the set of constant acts.

For any pair of events $B, E \subseteq \Omega$, $B \backslash E$ shall denote the set of elements that are in B but not in E. For any pair of acts f and $\widehat{f}$ in $\mathcal{F}$ and any event $B \subset \Omega$, $f_B \widehat{f}$ is identified with the act that agrees with f on B and with $\widehat{f}$ on $\Omega \backslash B$.

Let S denote the universe of conceivable realizations for any possible signal. It includes the *null* realization o.

The first component of an *information decision problem (IDP)* is a signal (structure) described by a function $\sigma \colon \Omega \to S$ with finite range, with the interpretation that the decision-maker learns the realization of the signal is s should the event $\sigma^{-1}(s)$ obtain. Let Σ denote the set of signals which I take to be a non-empty subset of functions from Ω to S. In particular, I assume Σ contains a *null* (that is, *uninformative*) signal, denoted $\mathbf{0}$, a constant mapping in which $\mathbf{0}(\omega) = o$ $(\in S)$ for all ω in Ω. The second component of an IDP is a finite menu of acts $M \subset \mathcal{F}$, $|M| < \infty$. Let $\mathcal{M}$ denote the set of menus. As any signal $\sigma \in \Sigma$ may be paired with any menu $M \in \mathcal{M}$ to form the IDP $(\sigma\,, M)$, the set of IDPs thus comprises the product set $\Sigma \times \mathcal{M}$.

A *plan (of action)* ϕ assigns to each $s \in S$ an act $\phi_s \in \mathcal{F}$. A plan of action ϕ is *feasible* in the IDP $(\sigma\,, M)$, if $\phi_s \in M$ for each $s \in \sigma(\Omega)$. With a slight abuse of the notation, each feasible plan in an IDP is *identified* with the act it generates. That is, a feasible plan ϕ for the IDP $(\sigma\,, M)$ is identified with the act defined by setting $\phi(\omega) := \phi_s(\omega)$ if $\omega \in \sigma^{-1}(s)$.

For each IDP $(\sigma\,, M)$ in $\Sigma \times \mathcal{M}$, let $\Phi(\sigma, M) \subset \mathcal{F}$ denote the set of acts generated by feasible plans in that IDP. I refer to any pair (σ, f) in $\Sigma \times \mathcal{F}$ as a *signal-act* and note it may be viewed as a the subclass of IDPs in which the menu contains just a single option. Thus I view and treat $\Sigma \times \mathcal{F}$ (the set of signal-acts) as a subset of $\Sigma \times \mathcal{M}$ (the set of IDPs).

The decision-maker is characterized by a binary relation $\succsim$ defined over $\Sigma \times \mathcal{F}$. For each particular signal σ in Σ, I refer to the restriction of $\succsim$ to $\{\sigma\} \times \mathcal{F}$ as his *ex ante preferences* with respect to that signal. Finally, I refer to any signal-act in $\{\mathbf{0}\} \times \mathcal{F}$ as *static* and the restriction of $\succsim$ to $\{\mathbf{0}\} \times \mathcal{F}$ as the decision-maker's *static preferences.*

A *prior* is a countably-additive, complete and non-atomic probability measure defined on a σ-algebra of subsets of Ω. Let Π denote the set of all priors with generic element π. For each π in Π (with domain $\mathcal{E}^\pi$) let $\mathcal{F}^\pi$ denote the set of acts that are measurable with respect to π (that is, adapted to $\mathcal{E}^\pi$).

Like his counterpart in Gul and Pesendorfer (2014), our decision-maker's static preferences are characterized by a pair $\langle \mu\,, u \rangle$, where:

(*i*) $\mu \in \Pi$ with domain $\mathcal{E}$ is the decision-maker's prior;[7] and,

(*ii*) $u\colon \{(x\,, y) \in X \times X \colon x \leqslant y\} \to \mathbb{R}$ is his interval utility, a continuous function for which $u(x\,, y) > u(x'\,, y')$ whenever $x > x'$ and $y > y'$.[8]

Adopting Gul and Pesendorfer's (2014) nomenclature, I will refer to any event E (respectively, any act g) in $\mathcal{E}$ (respectively, in $\mathcal{F}^\mu$) as *ideal.*

For any ideal act $g \in \mathcal{F}^\mu$ its static subjective expected utility is given by $\int v(g)\, d\mu$, where $v\colon X \to \mathbb{R}$ is the (increasing and continuous) function defined by setting $v(x) := u(x\,, x)$. For any arbitrary act f, however, although it may not be measurable with respect to μ, the decision-maker can always

[7] To ease the notational clutter, I let $\mathcal{E}$ rather than $\mathcal{E}^\mu$ denote the domain of the decision-maker's prior.

[8] A characterization of static EUU maximization can be found in Grant et al. (2024) which resolves some problems they identified in Gul and Pesendorfer's axiomatization.

compute ideal greatest-lower- and least-upper-bounds, that together form a (μ-)measurable *envelope* for that act.

Definition 1 (Greatest-lower- & least-upper-ideal bounds). *Fix a prior $\mu \in \Pi$. For each act $f \in \mathcal{F}$, the pair of acts $\underline{f}^{\mu}, \overline{f}^{\mu} \in \mathcal{F}^{\mu}$ constitutes an* ideal envelope *if $\underline{f}^{\mu} \leqslant f \leqslant \overline{f}^{\mu}$ and for any $\widehat{f} \in \mathcal{F}^{\mu}$:*

(i) $f \geqslant \widehat{f} \implies \mu\Big(\{\omega \in \Omega \colon \underline{f}^{\mu}(\omega) < \widehat{f}(\omega)\}\Big) = 0$;

(ii) $f \leqslant \widehat{f} \implies \mu\Big(\{\omega \in \Omega \colon \overline{f}^{\mu}(\omega) > \widehat{f}(\omega)\}\Big) = 0$;

Lemma 1 in Gul and Pesendorfer (2014, p5) establishes that for each act f there exists an essentially unique ideal envelope. The utility the decision-maker assigns the signal-act $(\mathbf{0}, f)$ is the (static) expected uncertain utility of f. That is,

$$V(\mathbf{0}, f) = \int u(\underline{f}^{\mu}, \overline{f}^{\mu})\, d\mu\,, \tag{6}$$

represents the restriction of $\succsim$ to $\{\mathbf{0}\} \times \mathcal{F}$ (his static preferences). Henceforth, I refer to such a decision-maker as an EUU maximizer.

Turning now to arbitrary signal-acts, recall from the introduction, I proposed that for each signal σ and each possible realization s of σ, the EUU maximizer attaches a posterior belief μ_s^{σ} defined over the events in the (relative) sigma-algebra

$$\mathcal{E}_s^{\sigma} = \{E \cap \sigma^{-1}(s) \colon E \in \mathcal{E}\}\,. \tag{7}$$

For each posterior μ_s^{σ}, let $\mathcal{F}^{\mu_s^{\sigma}} \subset \mathcal{F}$ denote the set of acts whose restriction to the information cell $\sigma^{-1}(s)$ is measurable with respect to μ_s^{σ}. Analogous to Definition 1 above of an ideal envelope, the next definition specifies a μ_s^{σ}-measurable envelope for the restriction of an act to the information cell $\sigma^{-1}(s)$.

Definition 2 (Greatest-lower- and least-upper- (μ_s^{σ}-)measurable bounds). *Fix a prior $\mu \in \Pi$, a signal $\sigma \in \Sigma$, and a realization $s \in \sigma(\Omega)$. For each posterior μ_s^{σ} (defined over $\mathcal{E}_s^{\sigma}$), and each act $f \in \mathcal{F}$, the pair of acts $\underline{f}^{\mu_s^{\sigma}}, \overline{f}^{\mu_s^{\sigma}} \in \mathcal{F}^{\mu_s^{\sigma}}$ that agree with f on $\Omega \backslash \sigma^{-1}(s)$, constitutes a μ_s^{σ}-*measurable envelope *if $\underline{f}^{\mu_s^{\sigma}} \leqslant f \leqslant \overline{f}^{\mu_s^{\sigma}}$ and for any $\widehat{f} \in \mathcal{F}^{\mu_s^{\sigma}}$ that agrees with f on $\Omega \backslash \sigma^{-1}(s)$:*

(i) $f \geqslant \widehat{f} \implies \mu_s^\sigma\Big(\{\omega \in \sigma^{-1}(s)\colon \underline{f}^{\mu_s^\sigma}(\omega) < \widehat{f}(\omega)\}\Big) = 0$;

(ii) $f \leqslant \widehat{f} \implies \mu_s^\sigma\Big(\{\omega \in \sigma^{-1}(s)\colon \overline{f}^{\mu_s^\sigma}(\omega) > \widehat{f}(\omega)\}\Big) = 0$;

To describe how I propose the decision-maker forms these posterior beliefs entails the use of the *ideal split* generated by the preimage of a signal. The specification of this partition of the state space requires some additional definitions and notation, starting with the inner-measure of μ, denoted by μ_* and defined by setting

$$\mu_*(B) := \sup_{E \subseteq B : E \in \mathcal{E}} \mu(E) \text{ for each } B \subseteq \Omega\,.$$

Since μ is countably additive, the infimum is attained. I shall refer to the ideal event $[B]_*$ in $\mathcal{E}$, as the *inner-sleeve* of B, if $[B]_* \subseteq B$ and $\mu([B]_*) = \mu_*(B)$. The inner-sleeve of B may be viewed as the largest ideal subset of B.

Definition 3 (A Signal's Ideal Split). *Fix a prior $\mu \in \Pi$. For each signal $\sigma \in \Sigma$, its* ideal split*, denoted by the collection of ideal events $\{E_Q^\sigma \in \mathcal{E} : Q \subseteq \sigma(\Omega)\,, Q \neq \emptyset\}$, is a partition of the state space inductively defined as follows:*

1. *For each realization $s \in \sigma(\Omega)$, set $E_{\{s\}}^\sigma := [\sigma^{-1}(s)]_*$.*
2. *For each $Q \subseteq \sigma(\Omega)$ such that $|Q| > 1$, set*

$$E_Q^\sigma := \left[\sigma^{-1}(Q)\right]_* \setminus \left(\bigcup_{\widehat{Q} \subset Q} E_{\widehat{Q}}^\sigma\right)\,.$$

*I refer to E_Q^σ as the σ-*marginal inner-sleeve *of the set of signal realizations Q.*

As an illustration, consider the ideal split generated by the preimage of a binary signal $\widehat{\sigma}$ that maps each state to one of two possible realizations, s' or s''. The ideal split is a three element partition of the state-space

$$\Big\{ \underset{\left[\widehat{\sigma}^{-1}(s')\right]_*}{\underset{\parallel}{E_{\{s'\}}^{\widehat{\sigma}}}}\,,\quad \underset{\left[\widehat{\sigma}^{-1}(s'')\right]_*}{\underset{\parallel}{E_{\{s''\}}^{\widehat{\sigma}}}}\,,\quad \underset{\Omega\setminus\left(\left[\widehat{\sigma}^{-1}(s')\right]_* \cup \left[\widehat{\sigma}^{-1}(s'')\right]_*\right)}{\underset{\parallel}{E_{\{s',s''\}}^{\widehat{\sigma}}}} \Big\}\,.$$

The first (respectively, second) element corresponds to the largest ideal subset in which the signal's realization is s' (respectively, s''). The third element corresponds to the $\widehat{\sigma}$-marginal inner-sleeve of the set of realizations $\widehat{\sigma}(\Omega)$ ($= \{s', s''\}$). All the EUU maximizer can discern from this element of the ideal split is that the realization will be either s' or s''. However, he is unable to attribute any fraction of the probability his prior assigns to this element of the split to either s' or s'' alone obtaining.

More generally for an arbitrary signal σ and for each non-empty set of its realizations $Q \subseteq \sigma(\Omega)$, recalling the approach of Dempster (1967) and Shafer (1976), I interpret $\mu\big(E_Q^\sigma\big)$ (the probability assigned by the prior to the σ-marginal inner-sleeve of Q) as measuring the weight the decision-maker places on the evidence that directly supports the realization of the signal residing in Q that *cannot be further refined* ex ante *in terms of any of its strict subsets.*

Since he anticipates after learning the realization of the signal he will then be assigning the information cell associated with that realization a (precise) posterior belief of 1 and all the other information cells (precise) posterior beliefs of 0, he treats the partition of information cells as *conditionally* ideal. Furthermore, given the particular realization s, he treats as *conditionally ideal* the intersection of each (unconditionally) ideal event with the information cell $\sigma^{-1}(s)$. That is, for each realization s of the signal, he treats as conditionally ideal all the events in the (relative) sigma-algebra $\mathcal{E}_s^\sigma$.

To help motivate the updating rule that appears in Definition 4 below, first notice for each non-empty set of the signal realizations $Q \subset \sigma(\Omega)$ and any pair of realizations s' and s'' that both reside in Q, we have

$$\Big(\mathbf{0}, \big(x_{\sigma^{-1}(s')}y\big)_{E\cap E_Q^\sigma} f\Big) \sim \Big(\mathbf{0}, \big(x_{\sigma^{-1}(s'')}y\big)_{E\cap E_Q^\sigma} f\Big) \ ,$$

for all ideal events $E \in \mathcal{E}$, all pairs of outcomes x and y, and all acts f. To see this, notice the application of the representation of the static preferences in equation (6) yields

$$\begin{aligned} &V\Big(\mathbf{0}, \big(x_{\sigma^{-1}(s')}y\big)_{E\cap E_Q^\sigma} f\Big) - V\Big(\mathbf{0}, \big(x_{\sigma^{-1}(s'')}y\big)_{E\cap E_Q^\sigma} f\Big) \\ &= \mu\left(E\cap E_Q^\sigma\right) u(\min\{x,y\}, \max\{x,y\}) + \int_{\Omega\setminus\left(E\cap E_Q^\sigma\right)} f\, d\mu \\ &\quad - \left[\mu\left(E\cap E_Q^\sigma\right) u(\min\{x,y\}, \max\{x,y\}) + \int_{\Omega\setminus\left(E\cap E_Q^\sigma\right)} f\, d\mu\right] = 0\,. \end{aligned}$$

This means from his static preferences we see the EUU maximizer is indifferent between betting on (respectively, against) any event in the collection

$\{E \cap E_Q^\sigma \cap \sigma^{-1}(s) \colon s \in Q\}$ versus betting on (respectively, against) any other event in the collection. In other words, all these events that together form a $|Q|$-element partition of the ideal event $E \cap E_Q^\sigma$ are *exchangeable* with each other. So although these events are not measurable with respect to the prior μ, the EUU maximizer's conditional belief system embodies a "principle of indifference" property by allocating the probability $\mu\big(E \cap E_Q^\sigma\big)$ *uniformly* across its intersections with the information cells associated with the realizations in Q.

The logic of this reasoning, underpins the next definition of the *uniform conditional belief system* that the EUU maximizer associates with a signal. It comprises a pseudo-prior (that renders all the information cells of the signal unambiguous) and in addition, for each realization for which the pseudo-prior is strictly positive, it specifies a conditional probability (his *posterior*) defined over those events he views as conditionally ideal.

Definition 4 (Uniform Conditional Belief System). *Fix a prior $\mu \in \Pi$. The* uniform conditional belief system *associated with the signal $\sigma \in \Sigma$ comprises for each realization $s \in \sigma(\Omega)$:*

(i) *a* pseudo-prior $p_s^\sigma \in [0\,,1]$ *given by*

$$p_s^\sigma := \sum_{Q \subseteq \sigma(\Omega) \colon s \in Q} \frac{\mu\left(E_Q^\sigma\right)}{|Q|}\,;$$

and, in addition, whenever $p_s^\sigma > 0$,

(ii) *a posterior μ_s^σ defined on $\mathcal{E}_s^\sigma = \{E \cap \sigma^{-1}(s) \colon E \in \mathcal{E}\}$ and obtained by setting for each event $E \in \mathcal{E}$:*

$$\mu_s^\sigma(E \cap \sigma^{-1}(s)) := \frac{\sum_{Q \subseteq \sigma(\Omega) \colon s \in Q} \mu\big(E \cap E_Q^\sigma\big)\,/|Q|}{p_s^\sigma}\,.$$

Remark 1. *Notice that $\mu_s^\sigma(E \cap \sigma^{-1}(s))$ is well-defined since for any pair of ideal events E and $\widehat{E}$ in $\mathcal{E}$, if $E \cap \sigma^{-1}(s) = \widehat{E} \cap \sigma^{-1}(s)$ then it readily follows that $\mu\big(E \cap E_Q^\sigma\big) = \mu\Big(\widehat{E} \cap E_Q^\sigma\Big)$ for all $Q \subseteq \sigma(\Omega)$ such that $s \in Q$.*

Remark 2. *Fix a signal $\sigma \in \Sigma$. If for any of its realizations $s \in \sigma(\Omega)$ the information set $\sigma^{-1}(s)$ is ideal, then as this means the information cell is equal to its inner-sleeve $[\sigma^{-1}(s)]_*$ it follows for every $Q \subseteq \sigma(\Omega)$ in which*

$s \in Q$ and $|Q| > 1$, the corresponding element in the ideal split E_Q^σ is empty. Hence we have

$$p_s^\sigma = \mu\big(\sigma^{-1}(s)\big)$$
$$\text{and } \ \mu_s^\sigma(E \cap \sigma^{-1}(s)) = \frac{\mu(E \cap \sigma^{-1}(s))}{\mu(\sigma^{-1}(s))}\,.$$

That is, the posterior corresponds to the standard Bayesian update of the prior conditional on that realization obtaining.

Remark 3. *The uniform conditional belief system exhibits the following* "law of total probability"*. For each $E \in \mathcal{E}$:*

$$\sum_{s \in S:\, p_s^\sigma > 0} p_s^\sigma \mu_s^\sigma\big(E \cap \sigma^{-1}(s)\big) = \mu(E)\,.$$

To see this, applying the definitions above we have

$$\sum_{s \in S:\, p_s^\sigma > 0} p_s^\sigma \mu_s^\sigma\big(E \cap \sigma^{-1}(s)\big) = \sum_{s \in S:\, p_s^\sigma > 0} \left(\sum_{Q \subseteq \sigma(\Omega):\, s \in Q} \frac{\mu\big(E \cap E_Q^\sigma\big)}{|Q|} \right)$$
$$= \sum_{Q \subseteq \sigma(\Omega)} \left(\sum_{s \in Q} \frac{\mu\big(E \cap E_Q^\sigma\big)}{|Q|} \right) = \sum_{Q \subseteq \sigma(\Omega)} \mu\big(E \cap E_Q^\sigma\big) = \mu(E)\ .$$

Furthermore, substituting Ω for E yields $\sum_s p_s^\sigma \underbrace{\mu_s^\sigma\big(\sigma^{-1}(s)\big)}_{=1} = \mu(\Omega) = 1$.

Example 2 (Pseudo-priors and posteriors for a non-ideal binary signal). *Set $\Omega := [0,1] \times [0,1]$ and let $\mathcal{B}$ denote the collection of Borel subsets of $[0,1]$. Take the set of ideal sets to be the σ-algebra constructed by setting*

$$\mathcal{E} := \{[0,1] \times B : \ B \in \mathcal{B}\},$$

so ideal events are formed from full-length horizontal bands. Let μ be the prior on $\mathcal{E}$ defined by

$$\mu([0,1] \times B) := \lambda(B),$$

where λ is Lebesgue measure on $[0,1]$.

Fix numbers $0 < a < b < 1$ and $\alpha \in (0,1)$. Define a binary signal σ with realizations $\{s', s''\}$ by specifying the information cell for s' as

$$\sigma^{-1}(s') = A := \big([0,1] \times [b,1]\big) \ \cup \ \big([0,\alpha] \times [a,b]\big),$$

and hence $\sigma^{-1}(s'') = \Omega \backslash A$.[9]

The ideal (measurable) split. *The largest $\mathcal{E}$-measurable subset of A is*

$$[A]_* = [0,1] \times [b,1],$$

and the largest $\mathcal{E}$-measurable subset of $\Omega \backslash A$ is

$$[\Omega \backslash A]_* = [0,1] \times [0,a].$$

Accordingly, the induced ideal split has three (ideal) cells:

$$E^\sigma_{\{s'\}} = [0,1] \times [b,1], \qquad E^\sigma_{\{s''\}} = [0,1] \times [0,a], \qquad E^\sigma_{\{s',s''\}} = [0,1] \times [a,b].$$

The third cell $E^\sigma_{\{s',s''\}}$ is the "residual" region: ex ante it can only be assigned to the coalition $\{s', s''\}$ (it is not ideal-refinable into singletons), even though ex post the signal realization determines whether the left or right block of it is reached.

Pseudo-priors (uniform split on the residual coalition). *Definition 4 assigns pseudo-priors*

$$p^\sigma_{s'} = \mu\big(E^\sigma_{\{s'\}}\big) + \frac{1}{2}\,\mu\big(E^\sigma_{\{s',s''\}}\big)\,, \qquad p^\sigma_{s''} = \mu\big(E^\sigma_{\{s''\}}\big) + \frac{1}{2}\,\mu\big(E^\sigma_{\{s',s''\}}\big)\,.$$

Since $\mu([0,1] \times [c,d]) = d - c$, we have

$$\mu\big(E^\sigma_{\{s'\}}\big) = 1 - b, \qquad \mu\big(E^\sigma_{\{s''\}}\big) = a, \qquad \mu\big(E^\sigma_{\{s',s''\}}\big) = b - a,$$

and therefore

$$p^\sigma_{s'} = (1-b) + \frac{1}{2}(b-a), \qquad p^\sigma_{s''} = a + \frac{1}{2}(b-a),$$

with $p^\sigma_{s'} + p^\sigma_{s''} = 1$.

[9] Geometrically: $\sigma^{-1}(s')$ is the top horizontal band $[0,1] \times [b,1]$ together with the left block of the middle band $[0,1] \times [a,b]$; its complement is the bottom band $[0,1] \times [0,a]$ together with the right block of the middle band.

Posteriors on conditionally ideal events. *Upon observing s', the posterior $\mu^\sigma_{s'}$ on $\mathcal{E}^\sigma_{s'} := \{E \cap A : \; E \in \mathcal{E}\}$ is defined by*

$$\mu^\sigma_{s'}(E \cap A) = \frac{\mu\Big(E \cap E^\sigma_{\{s'\}}\Big) + \frac{1}{2}\,\mu\Big(E \cap E^\sigma_{\{s',s''\}}\Big)}{p^\sigma_{s'}}, \qquad E \in \mathcal{E},$$

and similarly, upon observing s'',

$$\mu^\sigma_{s''}(E \cap (\Omega \backslash A)) = \frac{\mu\Big(E \cap E^\sigma_{\{s''\}}\Big) + \frac{1}{2}\,\mu\Big(E \cap E^\sigma_{\{s',s''\}}\Big)}{p^\sigma_{s''}}, \qquad E \in \mathcal{E}.$$

To see "where this bites" take the ideal test event $E_t = [0,1] \times [0,t]$ with $t \in [a,b]$. Then $E_t \cap E^\sigma_{\{s'\}} = \varnothing$ and $E_t \cap E^\sigma_{\{s',s''\}} = [0,1] \times [a,t]$, so

$$\mu^\sigma_{s'}(E_t \cap A) = \frac{\frac{1}{2}\,\mu([0,1] \times [a,t])}{(1-b) + \frac{1}{2}(b-a)} = \frac{\frac{1}{2}(t-a)}{(1-b) + \frac{1}{2}(b-a)}.$$

This makes explicit that probability mass from the residual (middle) band $[0,1] \times [a,b]$ is shared uniformly across $\{s', s''\}$, while singleton-refinable mass from either the top or bottom band is treated in the usual Bayesian way.

Relation to Dempster–Shafer style optimistic/pessimistic rules. In Dempster–Shafer updating, the residual mass attached to a coalition Q can be converted into lower/upper bounds (belief/plausibility) or, equivalently, allocated in an optimistic or pessimistic manner when forming point probabilities. REUU instead selects a uniform split across Q by an invariance/symmetry requirement: within any coalition Q that is not further ideal-refinable ex ante, the singletons in Q are treated as exchangeable, so no singleton is privileged ex ante. Comparative statics against "optimistic" or "pessimistic" allocations can be obtained by replacing the uniform split $\mu(E^\sigma_Q)/|Q|$ with an alternative allocation rule over Q; Example 2 can then be rerun verbatim to show exactly where such alternatives bite.

All the elements required to define the class of Recursive EUU preferences have now been assembled.

Definition 5 (Recursive EUU Preferences). *A binary relation $\succsim$ over signal-acts is a* Recursive Expected Uncertain Utility Preference *if there exists an interval utility u and a prior μ (with which we can associate for each* *signal $\sigma \in \Sigma$ the uniform conditional belief system $\left\langle (p_s^\sigma)_{s\in\sigma(\Omega)}, (\mu_s^\sigma)_{s:\, p_s^\sigma>0} \right\rangle$), such that $\succsim$ admits a representation of the form:*

$$V(\sigma, f) = \int u\Big(\underline{f_\sigma}^\mu, \overline{f}_\sigma^\mu\Big)\; d\mu\,,$$

where f_σ is an act adapted to σ that is defined by setting

$$f_\sigma(\omega) := \begin{cases} v^{-1}\circ V_s^\sigma(f) & \text{if } \omega \in \sigma^{-1}(s) \text{ and } p_s^\sigma > 0 \\ \ell & \text{otherwise} \end{cases}\;;$$

where v is the Bernoulli utility derived from u, defined by setting $v(x) := u(x,x)$, and for each realization s such that $p_s^\sigma > 0$, $V_s^\sigma \colon \mathcal{F} \to \mathbb{R}$ is defined by setting

$$V_s^\sigma(f) := \int_{\sigma^{-1}(s)} u\left(\underline{f}^{\mu_s^\sigma}, \overline{f}^{\mu_s^\sigma}\right)\, d\mu_s^\sigma\,.$$

3. Characterization

I begin by assuming the decision-maker's static preferences admit an EUU representation and that for each signal-act there exists a *certainty equivalent* signal-act, comprising the uninformative signal and a constant act (that is, outcome).[10]

Assumption 1 (Static EUU and Certainty Equivalent Existence). *The static preferences admit a representation of the form*

$$V(\mathbf{0}, f) = \int u\left(\underline{f}^\mu, \overline{f}^\mu\right)\, d\mu$$

for some prior $\mu \in \Pi$ (with domain $\mathcal{E}$) and interval utility u.

Moreover, for each signal-act $(\sigma\,, f) \in \Sigma \times \mathcal{F}$ there exists an outcome $x \in X$ such that $(\mathbf{0}\,, x) \sim (\sigma\,, f)$.

[10] An axiomatization of static EUU maximization can be found in Grant et al. (2024) which fixes a couple of problems they identified in the original characterization result in Gul and Pesendorfer (2014). Standard ordering and continuity axioms are sufficient to ensure the existence of a certainty equivalent for each signal-act.

Notice it follows from the monotonicity of the interval utility u that the certainty equivalent of any signal-act is well defined.

Recall in the introduction, I noted that Gul and Pesendorfer viewed the set of ideal events (the domain of the prior μ) as ones capturing aspects of the uncertainty that the decision maker could quantity precisely with a probability. In terms of the static preferences this requires for any ideal event $E \in \mathcal{E}$ that both it and its complement exhibit a form of Savage's property **P2**.

Definition 6 (Ideal Event). *The event $E \subseteq \Omega$ is* ideal *if for every quartet of acts f, f', $\widehat{f}$ and $\widetilde{f}$ in $\mathcal{F}$:*

$$\left[(\mathbf{0}, f_E\widehat{f}) \succsim (\mathbf{0}, f'_E\widehat{f}) \text{ and } (\mathbf{0}, \widehat{f}_E f) \succsim (\mathbf{0}, \widehat{f}_E f')\right]$$
$$\Longrightarrow$$
$$\left[(\mathbf{0}, f_E\widetilde{f}) \succsim (\mathbf{0}, f'_E\widetilde{f}) \text{ and } (\sigma, \widetilde{f}_E f) \succsim (\mathbf{0}, \widetilde{f}_E f')\right] .$$

To motivate the first axiom, notice that any signal-act in which the act is adapted to the signal has all uncertainty resolved in one stage: once the decision-maker observes the signal realization, the resulting outcome is known, so uncertainty resolves "early." By contrast, a signal-act involving the uninformative signal $\mathbf{0}$ also resolves all uncertainty in one stage, but only "late." Axiom 1 therefore imposes time neutrality: the REUU maximizer is indifferent between these two timing descriptions of the same adapted act. Unlike the compound risk models of Kreps and Porteus (1978), Epstein and Zin (1989) and Grant et al. (1998), REUU therefore rules out an intrinsic strict preference for early rather than late resolution of uncertainty.

Axiom 1 (Time Neutrality). *For each signal $\sigma \in \Sigma$: $(\mathbf{0}\,, f) \sim (\sigma\,, f)$ for every act in which $\sigma(\omega) = \sigma(\widehat{\omega}) \implies f(\omega) = f(\widehat{\omega})$.*

Next, I require the conditional ex ante preferences associated with a given signal and any of its realizations should not depend on how the acts behave on states outside the information cell corresponding to that realization. Formally, this entails the ex ante preferences with respect to a given signal are weakly separable across the partition of information cells generated by the signal's realizations.

Axiom 2 (Within-Event Consequentialism). *Fix a signal $\sigma \in \Sigma$. For all realizations $s \in \sigma(\Omega)$ and all pairs of acts $f, g \in \mathcal{F}$: $(\sigma\,, f) \succsim \left(\sigma\,, g_{\sigma^{-1}(s)}f\right) \implies \left(\sigma\,, f_{\sigma^{-1}(s)}g\right) \succsim (\sigma\,, g)$.*

In addition, for a given signal, I require the intersection of each ideal event in $\mathcal{E}$ with each information cell of the signal to be "conditionally ideal". Formally this entails, relative to each information cell, each ideal event in $\mathcal{E}$ and its complement jointly satisfy Savage's postulate **P2**.[11]

Axiom 3 (Ideal Invariance). *Fix a signal $\sigma \in \Sigma$. For each realization $s \in \sigma(\Omega)$, each (ideal) event $E \in \mathcal{E}$, and each pair of acts $f, g \in \mathcal{F}$:*

$$\big(\sigma\,, f_{\sigma^{-1}(s)}f\big) \succsim \big(\sigma\,, (g_E f)_{\sigma^{-1}(s)}f\big) \implies \big(\sigma\,, (f_E g)_{\sigma^{-1}(s)}f\big) \succsim \big(\sigma\,, g_{\sigma^{-1}(s)}f\big)$$

and

$$\big(\sigma\,, f_{\sigma^{-1}(s)}f\big) \succsim \big(\sigma\,, (f_E g)_{\sigma^{-1}(s)}f\big) \implies \big(\sigma\,, (g_E f)_{\sigma^{-1}(s)}f\big) \succsim \big(\sigma\,, g_{\sigma^{-1}(s)}f\big)$$

Up to this point, the axioms pin down (i) the ex ante evaluation of acts on the ideal σ-algebra $\mathcal{E}$ (via the static EUU representation) and (ii) the way interim preferences are derived from the ex ante ranking over signal-acts. What remains open is how the agent evaluates *uncertain outcomes* generated within a realized information cell when that cell is not itself ideal. Absent an additional restriction, one could rationalize interim behavior using interval utilities that vary with the information cell or the signal. Axiom 4 rules out such state-dependent reparametrizations by requiring that the same interval utility u that represents static EUU also governs the evaluation of uncertain outcomes at every history. In short, Axiom 4 is the invariance condition that delivers *model consistency* of the interim criterion with EUU: it ensures that information affects evaluation only through the (posterior) belief component, not through a change in the agent's attitude to risk and ambiguity. To state it requires the following additional definitions and notation.

An event $B \subset \Omega$ is *null (with respect to the static preferences)* if $(\mathbf{0}, g_B f) \sim (\mathbf{0}, f)$ for all $f, g \in \mathcal{F}$. Let $\mathcal{N}$ denote the set of null events.[12]

For each non-null event $B \notin \mathcal{N}$ and each act $f \in \mathcal{F}$, set

$$\underline{x}(f_B) := \sup\{y \in X\colon \{\omega \in B\colon f(\omega) < y\} \in \mathcal{N}\}\,.$$
$$\overline{x}(f_B) := \inf\{z \in X\colon \{\omega \in B\colon f(\omega) > z\} \in \mathcal{N}\}\,.$$

---

[11] That is, relative to each information cell, each event E in $\mathcal{E}$ is both what Gul and Pesendorfer (2014) refer to as "left-ideal" and "right-ideal". It then follows from Lemma B0 in Gul and Pesendorfer (2014, p25) that relative to each information cell of the signal such an event satisfies definition 6.

[12] An event is null if and only if μ assigns it measure zero.

In terms of the EUU maximizer's static preferences, I interpret $\underline{x}(f_B)$ (respectively, $\overline{x}(f_B)$) as the greatest lower (respectively, least upper) bound for the outcomes the EUU maximizer conceives as possibly arising on the restriction of f to the event B. I write $\underline{x}(f)$ for $\underline{x}(f_\Omega)$ and $\overline{x}(f)$ for $\overline{x}(f_\Omega)$.

For each signal $\sigma \in \Sigma$ set

$$\sigma(\Omega)^+ := \{s \in \sigma(\Omega) \colon \sigma^{-1}(s) \notin \mathcal{N}\} \,.$$

That is, $\sigma(\Omega)^+$ comprises all of that signal's non-null realizations.

An event D is deemed *diffuse (with respect to the static preferences)* if $D \cap E \neq \emptyset \neq (\Omega \backslash D) \cap E$ for every non-null ideal event E. A diffuse event may be viewed as maximally uncertain since neither it nor its complement contains a non-null ideal subset. Let $\mathcal{D}$ denote the set of diffuse events.

An act h is deemed *diffuse* if for some diffuse event $D \in \mathcal{D}$, $(\mathbf{0}, h) \sim (\mathbf{0}, \overline{x}(h)_D \underline{x}(h))$. Let $\mathcal{H}$ denote the set of diffuse acts. Notice, for any diffuse act $h \in \mathcal{H}$, its static EUU $V(\mathbf{0}, h) = V(\mathbf{0}, \overline{x}(h)_D \underline{x}(h)) = u(\underline{x}(h), \overline{x}(h))$.

Axiom 4 (Invariant Uncertain Outcome Preference). *Fix a signal $\sigma \in \Sigma$. For any non-null realization $s \in \sigma(\Omega)^+$, any act f in $\mathcal{F}$, any pair of diffuse acts h and $\widehat{h}$ in $\mathcal{H}$, any outcome z in X, and any ideal event $E \in \mathcal{E}$ for which $E \cap \sigma^{-1}(s) \notin \mathcal{N}$: if $\underline{x}(h) = \underline{x}(\widehat{h}_{E\cap\sigma^{-1}(s)})$ and $\overline{x}(h) = \overline{x}(\widehat{h}_{E\cap\sigma^{-1}(s)})$ then*

$$(\mathbf{0}\,, h) \succsim (\mathbf{0}\,, z) \iff \left(\sigma, \widehat{h}_{E\cap\sigma^{-1}(s)} f\right) \succsim \left(\sigma, z_{E\cap\sigma^{-1}(s)} f\right) \,.$$

Axiom 4 fixes only the utility side of interim evaluation. Information may change the agent's beliefs, but it may not change the interval utility u that represents his attitude toward risk and ambiguity. What remains open is which conditional belief over ideal events should govern interim rankings when the realized information cell $\sigma^{-1}(s)$ is not itself ideal. Axiom 5 resolves that remaining indeterminacy by fixing the belief-updating component: conditional rankings must agree with the ex ante rankings of the corresponding bets on the uniformly assigned pieces generated by the uniform refinement construction.

To motivate this updating rule, fix a signal $\sigma \in \Sigma$ and consider its associated ideal split $\{E^\sigma_Q \in \mathcal{E} \,:\, Q \subseteq \sigma(\Omega),\; Q \neq \emptyset\}$. Recall that for any $Q \subseteq \sigma(\Omega)$ and any ideal event $E \in \mathcal{E}$ with $E \cap E^\sigma_Q \notin \mathcal{N}$, the family

$$\left\{\, E \cap E^\sigma_Q \cap \sigma^{-1}(s) \;:\; s \in Q \,\right\}$$

forms a $|Q|$-element partition of the ideal event $E \cap E_Q^\sigma$. Moreover, because $\succsim$ evaluates any pair $(\mathbf{0}, f)$ through the $\mathcal{E}$-measurable bounds of f, the REUU maximizer views the elements of this partition as ex ante exchangeable: for any $s', s'' \in Q$, any $x > y$, and any act f, he is indifferent between betting on (or against) $E \cap E_Q^\sigma \cap \sigma^{-1}(s')$ and betting on (or against) $E \cap E_Q^\sigma \cap \sigma^{-1}(s'')$. Intuitively, on the residual ideal cell E_Q^σ the signal cannot be further ideal-refined ex ante, so the singletons $s \in Q$ are treated symmetrically.

To respect this exchangeability in the construction of realization-contingent beliefs, the REUU maximizer allocates the prior weight $\mu(E \cap E_Q^\sigma)$ *uniformly* across the $|Q|$ pieces $\{E \cap E_Q^\sigma \cap \sigma^{-1}(s) : s \in Q\}$. The next definition introduces the corresponding uniform refinement partition.

Definition 7 (Uniform refinement ideal partition). *Fix an ideal event $E \in \mathcal{E}$ and a signal $\sigma \in \Sigma$, with associated ideal split $\{E_Q^\sigma \in \mathcal{E} \,:\, Q \subseteq \sigma(\Omega),\ Q \neq \emptyset\}$. For each non-empty $Q \subseteq \sigma(\Omega)$ choose an ideal partition $\{E_{Q,s}^{\sigma,E} : s \in Q\}$ of E_Q^σ such that, for every $s \in Q$,*

$$\mu\Big(E_{Q,s}^{\sigma,E}\Big) = \frac{\mu(E_Q^\sigma)}{|Q|} \qquad \textit{and} \qquad \mu\Big(E_{Q,s}^{\sigma,E} \cap E\Big) = \frac{\mu(E_Q^\sigma \cap E)}{|Q|}.$$

Define the induced uniform refinement partition $\{E_s : s \in \sigma(\Omega)\}$ by

$$E_s \;:=\; \bigcup_{\substack{Q \subseteq \sigma(\Omega) \\ s \in Q}} E_{Q,s}^{\sigma,E}\,, \qquad s \in \sigma(\Omega).$$

By construction,

$$\mu(E \cap E_s) = \sum_{\substack{Q \subseteq \sigma(\Omega) \\ s \in Q}} \frac{\mu(E_Q^\sigma \cap E)}{|Q|} \qquad \text{and} \qquad \mu(E_s) = \sum_{\substack{Q \subseteq \sigma(\Omega) \\ s \in Q}} \frac{\mu(E_Q^\sigma)}{|Q|},$$

and the latter does not depend on E.

The last axiom formalizes this updating rule.

Axiom 5 (Invariant Conditional Betting Preferences). *Fix $E, E' \in \mathcal{E}$ and a signal $\sigma \in \Sigma$ with associated ideal split $\{E_Q^\sigma \in \mathcal{E} \,:\, Q \subseteq \sigma(\Omega),\ Q \neq \emptyset\}$. Let $\{E_s : s \in \sigma(\Omega)\}$ (respectively, $\{E'_s : s \in \sigma(\Omega)\}$) be a uniform refinement ideal partition associated with (E, σ) (respectively, (E', σ)). For each $s \in \sigma(\Omega)^+$,*

$$(\sigma, m_{E \cap \sigma^{-1}(s)}\ell) \;\succsim\; (\sigma, m_{E' \cap \sigma^{-1}(s)}\ell) \iff (\mathbf{0}, m_{E \cap E_s}\ell) \;\succsim\; (\mathbf{0}, m_{E' \cap E'_s}\ell)\,.$$

Thus Axiom 5 identifies the admissible posterior beliefs when the realized information cell is non-ideal: conditional rankings must respect the ex ante exchangeability encoded by the uniform refinement construction.

The representation theorem follows.

Theorem 1. *Suppose Assumption 1 holds. Then the binary relation $\succsim$ is an REUU preference if and only if Axioms 1 – 5 hold.*

4. Neo-Additive Signal-Acts

This section identifies a tractable subclass of signal-acts—*neo-additive signal-acts*—for which both ex ante and interim evaluations admit a parsimonious "neo-EU" representation. The representation parallels the neo-additive capacity model of Chateauneuf et al. (2007), but in REUU the interval utility u is unrestricted. When u is nonseparable, the resulting ambiguity term can generate Machina-style reflection patterns (violations of tail separability) that are impossible under any CEU model, including neo-additive CEU.

Neo-additive acts and signal-acts

Definition 8 (Neo-additive acts). *The act $g \in \mathcal{F}$ is* neo-additive with respect to the ideal event $G \in \mathcal{E}$ *if for all $x \geqslant y$ such that* $\mu_*(\{\omega \in \Omega : y \leq g(\omega) \leq x\}) \in (0,1)$,

1. $\big[\{\omega \in \Omega : y \leq g(\omega) \leq x\}\big]_* = \{\omega \in \Omega : y \leq g(\omega) \leq x\} \cap G$; *and*
2. $\Omega \backslash \big[\Omega \backslash \{\omega \in \Omega : y \leq g(\omega) \leq x\}\big]_* = \{\omega \in \Omega : y \leq g(\omega) \leq x\} \cup (\Omega \backslash G)$.

An act is neo-additive *if it is neo-additive with respect to some $G \in \mathcal{E}$.*

Neo-additivity with respect to G means: on the "ideal" region G the act behaves like an ideal act, while on the complement $\Omega \backslash G$ it is evaluated only through its worst and best conceivable prizes $\underline{x}(g)$ and $\overline{x}(g)$. In particular, the greatest lower and least upper *ideal* bounds of g are $g_G\underline{x}(g)$ and $g_G\overline{x}(g)$.

Definition 9 (Neo-additive signal-acts). *A signal-act $(\tau, g) \in \Sigma \times \mathcal{F}$ is* neo-additive *if there exists $G \in \mathcal{E}$ such that:*

1. *g is neo-additive with respect to G; and*

2. *the ideal split of τ satisfies*

$$G \backslash \Big(\bigcup_{s \in \tau(\Omega)^+} E^\tau_{\{s\}} \Big) \in \mathcal{N} \quad \textit{and} \quad (\Omega \backslash G) \backslash E^\tau_{\tau(\Omega)^+} \in \mathcal{N}.$$

For the null signal $\mathbf{0}$, $\mathbf{0}(\Omega)^+ = \{o\}$ and $E^{\mathbf{0}}_{\{o\}} = E^{\mathbf{0}}_{\mathbf{0}(\Omega)^+} = \Omega$, so condition 2 is automatic; hence $(\mathbf{0}, g)$ is neo-additive whenever g is neo-additive. When $|\tau(\Omega)^+| > 1$, condition 2 says (up to null events) that the singleton elements of the ideal split account for G, while the residual element $E^\tau_{\tau(\Omega)^+}$ accounts for $\Omega \backslash G$; in particular, $E^\tau_Q = \emptyset$ for all $Q \subset \tau(\Omega)^+$ with $|Q| > 1$.

Evaluating neo-additive signal-acts

Fix g neo-additive with respect to $G \in \mathcal{E}$. Writing $v(x) := u(x, x)$ and setting $\lambda := 1 - \mu(G)$, the static EUU of $(\mathbf{0}, g)$ is

$$\begin{aligned} V(\mathbf{0}, g) &= \int u\big(\underline{g}^\mu, \overline{g}^\mu\big) \, d\mu \\ &= \mu(G) \int v(g) \, d\pi + (1 - \mu(G)) \, u(\underline{x}(g), \overline{x}(g)) \\ &= (1 - \lambda) \int_X v \, dF^\pi_g + \lambda \, u(\underline{x}(g), \overline{x}(g)) \,, \end{aligned} \tag{8}$$

where (when $\mu(G) > 0$) π is the conditional probability on G,

$$\pi(B) := \frac{\mu(B \cap G)}{\mu(G)} \qquad \text{for } B \in \mathcal{E}^\pi := \{B : B \cap G \in \mathcal{E}\}, \tag{9}$$

and $F^\pi_g(x) := \pi(\{\omega : g(\omega) \le x\})$ is the π-CDF of g. In the separable special case $u(x, y) = \alpha v(x) + (1 - \alpha) v(y)$, (8) reduces to the Chateauneuf et al. neo-additive CEU formula; then the "tail increment" $u(a, m) - u(a, x) = (1 - \alpha)\big(v(m) - v(x)\big)$ is independent of the lower bound a. With nonseparable u, this independence can fail.

Example 3 (Neo-additive tail nonseparability). *Pick prizes $\ell < z < x < m$. Fix $G \in \mathcal{E}$ with $\mu(G) = 1 - \lambda \in (0, 1)$ and define π as in* (9)*. Pick an ideal event $E \in \mathcal{E}$ with $E \subseteq G$ and $\pi(E) = \frac{1}{2}$. Fix a diffuse event $D \in \mathcal{D}$ such that both $(\Omega \backslash G) \cap D$ and $(\Omega \backslash G) \cap (\Omega \backslash D)$ are nonnull.*

Define four neo-additive acts $f, g, \widetilde{f}, \widetilde{g}$ by their restrictions to G and $\Omega\backslash G$:

$$f|_G = \begin{cases} x, & \omega \in E, \\ z, & \omega \in G\backslash E, \end{cases} \qquad g|_G \equiv z, \qquad \widetilde{f}|_G \equiv z, \qquad \widetilde{g}|_G = \begin{cases} x, & \omega \in E, \\ z, & \omega \in G\backslash E, \end{cases}$$

and on the diffuse part

$$f|_{\Omega\backslash G} = \begin{cases} \ell, & \omega \in (\Omega\backslash G) \cap D, \\ m, & \omega \in (\Omega\backslash G) \cap (\Omega\backslash D), \end{cases} \qquad g|_{\Omega\backslash G} = \begin{cases} \ell, & \omega \in (\Omega\backslash G) \cap D, \\ x, & \omega \in (\Omega\backslash G) \cap (\Omega\backslash D), \end{cases}$$

$$\widetilde{f}|_{\Omega\backslash G} = \begin{cases} z, & \omega \in (\Omega\backslash G) \cap D, \\ m, & \omega \in (\Omega\backslash G) \cap (\Omega\backslash D), \end{cases} \qquad \widetilde{g}|_{\Omega\backslash G} = \begin{cases} z, & \omega \in (\Omega\backslash G) \cap D, \\ x, & \omega \in (\Omega\backslash G) \cap (\Omega\backslash D). \end{cases}$$

Then

$$(\underline{x}(f), \overline{x}(f)) = (\ell, m), \quad (\underline{x}(g), \overline{x}(g)) = (\ell, x),$$
$$(\underline{x}(\widetilde{f}), \overline{x}(\widetilde{f})) = (z, m), \quad (\underline{x}(\widetilde{g}), \overline{x}(\widetilde{g})) = (z, x).$$

Let $\Delta := \frac{1-\lambda}{2}\big(v(x) - v(z)\big) > 0$. *Using* (8) *and cancelling common terms,*

$$V(\mathbf{0}, f) - V(\mathbf{0}, g) = \Delta + \lambda\big(u(\ell, m) - u(\ell, x)\big),$$

$$V(\mathbf{0}, \widetilde{f}) - V(\mathbf{0}, \widetilde{g}) = -\Delta + \lambda\big(u(z, m) - u(z, x)\big).$$

Hence one obtains the reflection pattern $(\mathbf{0}, f) \succ (\mathbf{0}, g)$ *and* $(\mathbf{0}, \widetilde{g}) \succ (\mathbf{0}, \widetilde{f})$ *whenever*

$$u(\ell, m) - u(\ell, x) \ > \ \Delta/\lambda \ > \ u(z, m) - u(z, x). \tag{10}$$

Condition (10) *is incompatible with the separable neo-additive CEU benchmark, in which the increment* $u(a, m) - u(a, x)$ *is independent of* a.

Now consider a neo-additive signal-act (τ, g) with $|\tau(\Omega)^+| > 1$. For each $s \in \tau(\Omega)^+$, the pseudo-prior and posterior (Definition 4) simplify to

$$p_s^\tau = \mu(G \cap \tau^{-1}(s)) + \frac{\mu(\Omega\backslash G)}{|\tau(\Omega)^+|} = (1-\lambda)\pi(\tau^{-1}(s)) + \lambda\frac{1}{|\tau(\Omega)^+|},$$

$$\mu_s^\tau(E \cap \tau^{-1}(s) \cap G) = \frac{\mu(E \cap \tau^{-1}(s) \cap G)}{p_s^\tau} = \big(1 - \lambda_s^\tau\big)\, \pi_s^\tau(E) \qquad (E \in \mathcal{E}^\pi),$$

where π_s^τ is the Bayesian update of π on $\tau^{-1}(s)$,

$$\pi_s^\tau(B) := \frac{\pi(B \cap \tau^{-1}(s))}{\pi(\tau^{-1}(s))} \qquad (B \in \mathcal{E}^\pi),$$

and the updated lack-of-confidence parameter is

$$\lambda_s^\tau := \frac{\lambda/|\tau(\Omega)^+|}{(1-\lambda)\pi(\tau^{-1}(s)) + \lambda/|\tau(\Omega)^+|}. \tag{11}$$

Consequently, the interim preferences $\succsim_s^\tau$ admit a conditional neo-EU representation:

$$V_s^\tau(g) = (1-\lambda_s^\tau)\int_X v\, dF_g^{\pi_s^\tau} + \lambda_s^\tau\, u\big(\underline{x}(g_{\tau^{-1}(s)}), \overline{x}(g_{\tau^{-1}(s)})\big)\,, \tag{12}$$

where $F_g^{\pi_s^\tau}(x) := \pi_s^\tau(\{\omega : g(\omega) \leqslant x\})$.

Remark 4. *From* (11)*: (i) if* $|\tau(\Omega)^+| = 1$ *(an uninformative signal), then* $\lambda_s^\tau = \lambda$*; (ii) if* $|\tau(\Omega)^+| > 1$*, then* $\lambda_s^\tau < \lambda$ *iff* $\pi(\tau^{-1}(s)) > 1/|\tau(\Omega)^+|$ *(and conversely); (iii) if* $\pi(\tau^{-1}(s)) = 0$*, then* $\lambda_s^\tau = 1$ *(all confidence in* π *is lost on an event* π *deems impossible).*

Neo-additive sources

Definitions 8–9 characterize when a given act or signal-act is neo-additive. Since every ideal act is neo-additive (take $G = \Omega$), it is natural to ask whether there are priors π for which *all* π-measurable acts and $\mathcal{E}^\pi$-adapted signal-acts inherit a neo-EU representation.

Definition 10 (Neo-additive source). *A prior* $\pi \in \Pi$ *(with domain* $\mathcal{E}^\pi$*) is a* neo-additive source *if there exists a non-null* $G \in \mathcal{E}\backslash\mathcal{N}$ *such that*

$$\mu_*(A) = \begin{cases} \mu(G)\,\pi(A), & \text{if } \pi(A) < 1, \\ 1, & \text{if } \pi(A) = 1, \end{cases} \qquad \text{for all } A \in \mathcal{E}^\pi.$$

The term "source" is motivated by the fact that for all $A, B \in \mathcal{E}^\pi$ and all $x > y$,

$$V(\mathbf{0}, x_A y) \geqslant V(\mathbf{0}, x_B y) \iff \pi(A) \geqslant \pi(B),$$

matching the defining feature of sources in Gul and Pesendorfer (2015).

For any CDF F write

$$\underline{x}(F) := \sup\{y \in X : F(y) = 0\}, \qquad \overline{x}(F) := \inf\{z \in X : F(z) = 1\}.$$

Suppose the prior π is a neo-additive source, then for any $g \in \mathcal{F}^\pi$,

$$V(\mathbf{0}, g) = (1-\lambda) \int_X v \, dF_g^\pi + \lambda \, u\big(\underline{x}(F_g^\pi), \overline{x}(F_g^\pi)\big) , \tag{13}$$

with $\lambda = 1 - \mu(G)$. Moreover, if (τ, g) is such that both τ and g are adapted to $\mathcal{E}^\pi$, then for each $s \in \tau(\Omega)^+$,

$$V_s^\tau(g) = (1-\lambda_s^\tau) \int_X v \, dF_g^{\pi_s^\tau} + \lambda_s^\tau \, u\big(\underline{x}(F_g^{\pi_s^\tau}), \overline{x}(F_g^{\pi_s^\tau})\big) , \tag{14}$$

where π_s^τ is the Bayesian update of π on $\tau^{-1}(s)$ and λ_s^τ is given by (11). Hence both the static *and* the interim preferences are probabilistically sophisticated in the sense of Machina and Schmeidler (1992).

The next example gives a concrete interpretation of definition 10: G is the region on which probabilities are trusted (evaluation proceeds via the induced π-distribution), whereas $\Omega \backslash G$ contributes only through the extreme bounds term $u(\underline{x}, \overline{x})$.

Example 4 (A canonical neo-additive source). *Take Ω to be the unit square $[0,1]^2$ and let ν denote the Lebesgue measure on Ω. Set $\mathcal{E}$ equal to the σ-algebra of sets that are measurable in the second coordinate (up to ν-null exceptions in the first), and define the ideal prior μ on $\mathcal{E}$ by setting*

$$\mu(E) := \nu(E) \qquad \textit{for all } E \in \mathcal{E}.$$

Fix $\bar{y} \in (0,1)$ and set $G = [0,1] \times [0, \bar{y}] \in \mathcal{E}$. Define a prior π on the domain $\mathcal{E}^\pi := \{B \subseteq \Omega : \ B \cap G \in \mathcal{E}\}$ by

$$\pi(B) := \frac{\mu(B \cap G)}{\mu(G)} = \frac{\nu(B \cap G)}{\nu(G)} \qquad \textit{for all } B \in \mathcal{E}^\pi.$$

Then π is a neo-additive source.

Example 4 establishes neo-additive sources are not empty curiosities. Proposition 2 strengthens this by showing that, for every non-null ideal event G, one can construct a source that treats G as the "trusted" region and collapses the complement into the ambiguity term. Thus the neo-EU formulas above are supported by a broad existence result, not by a single canonical construction.

Proposition 2. *Fix a non-null event $G \in \mathcal{E} \backslash \mathcal{N}$. There exists a neo-additive source $\pi \in \Pi$ with domain $\mathcal{E}^\pi$ satisfying*

$$\pi(A) = \frac{\mu(A \cap G)}{\mu(G)} \qquad \textit{for each } A \in \mathcal{E}^\pi.$$

5. Related Literature.

For dynamic decision problems with non-additive beliefs and non-neutral attitudes toward ambiguity, two broad strands of work have developed.

A first, largely *statistical* strand studies updating rules for non-additive objects (capacities, sets of priors, belief functions) and the properties of the induced conditional expectations. Representative contributions include Denneberg (1994, 2002), Jaffray (1992), Lehrer (2005), Walley (1991), Shafer (1976), Lapied et al. (2012), and more recently Gul and Pesendorfer (2021) and Cerreia-Vioglio et al. (2024).

A second, *decision-theoretic* strand derives dynamic representations and updating rules from axioms imposed on ex ante and interim preferences; non-exhaustive examples include Gilboa and Schmeidler (1993), Sarin and Wakker (1998), Ghirardato (2002), Pires (2002), Epstein and Schneider (2003), Wang (2003), Hanany and Klibanoff (2007), Eichberger et al. (2007), and Siniscalchi (2011).

The present paper follows the decision-theoretic approach. It extends Gul and Pesendorfer's (2014) static expected uncertain utility (EUU) to *information decision problems* by introducing a conditional belief system (Definition 4) that specifies pseudo-priors and posteriors for *arbitrary* signals. The central feature is that, after any realized signal, the induced interim preferences retain an EUU form with the *same* interval utility u (model consistency), are consequentialist in the within-event sense used here, and yield dynamically consistent contingent plans. This contrasts with approaches that achieve dynamic consistency by relaxing consequentialism (e.g. Sarin and Wakker (1998), Siniscalchi (2011)) or with multiple-prior models that retain both consequentialism and dynamic consistency only under structural restrictions such as rectangularity (e.g. Epstein and Schneider (2003)).

The *interpretation* of Definition 4 is related in spirit to "extended Bayesianism" (Piermont, 2021): when a signal is not adapted to ideal events, learning its realization changes the agent's ability to refine uncertainty *within* the realized information cell. Although the agent's awareness does not expand, the distinction between ideal (quantifiable) and diffuse (non-quantifiable) components is effectively sharpened upon learning. The conditional belief system then specifies how the prior ideal content is mapped into interim beliefs over the relative σ-algebra associated with the realized cell.

A natural benchmark for comparison is Ghirardato (2002), who shows that combining certain notions of consequentialism and dynamic consistency can force ex ante preferences back to subjective expected utility. The reconciliation is that the interim preference family induced here is *signal-dependent*: even for a fixed realization s, the interim functional $V_s^\sigma(\cdot)$ may depend on the entire signal σ and not merely on the conditioning event $\sigma^{-1}(s)$. In this sense the present framework retains within-event consequentialism while eschewing the stronger signal-independence implicit in some tree-based formulations.

Finally, Gul and Pesendorfer (2021) develop a recursive and consequentialist theory of belief revision for two-stage uncertainty under non-additive beliefs. Their approach constructs, for each signal, a "proxy" object that renders the signal's cells unambiguous before applying Bayes-type conditioning within each cell. The present paper is closest in spirit to this proxy-then-condition logic, but differs in primitives and scope: the starting point here is EUU preferences represented by a convex-ranged prior μ and interval utility u, and the conditional belief system is derived so as to preserve EUU model consistency and the consequentialist/dynamic-consistency package on unrestricted signal/act domains. In applications, this delivers a tractable dynamic EUU extension without imposing a priori domain restrictions or rectangularity-type constraints.

Appendix A. Proofs

Proof of Theorem 1

Sufficiency of the axioms

By Assumption 1 there exists a prior $\mu \in \Pi$ with domain $\mathcal{E}$ and an interval utility u such that the static preferences are represented by the function $V(\mathbf{0}, f) = \int u\left(\underline{f}^{\mu}, \overline{f}^{\mu}\right) d\mu$. Without loss of generality, set $u(m,m) := 1$ and $u(\ell,\ell) := 0$.

Fix a signal $\sigma \in \Sigma$. Let

(*i*) $\{E_Q^\sigma \in \mathcal{E} : Q \subseteq \sigma(\Omega), Q \neq \emptyset\}$ denote the ideal split associated with σ as specified in definition 3; and, (*ii*) $\left\langle (p_s^\sigma)_{s\in\sigma(\Omega)}, (\mu_s^\sigma)_{s\colon p_s^\sigma>0} \right\rangle$ denote the uniform conditional belief system associated with the signal σ as specified in definition 4.

For each realization $s \in \sigma(\Omega)^+$, let $\succsim_s^\sigma$ denote the binary relation over acts derived from $\succsim$ by setting for each pair of acts f and g in $\mathcal{F}$, $f \succsim_s^\sigma g$ if $(\sigma\,, f) \succsim (\sigma\,, g_{\sigma^{-1}(s)} f)$.

From Axiom 2 (consequentialism) it readily follows that $f \succsim_s^\sigma g$ if and only if $f_{\sigma^{-1}(s)} h \succsim_s^\sigma g_{\sigma^{-1}(s)} h$ for *all* h in $\mathcal{F}$.[13]

Applying Assumption 1, let $c_s^\sigma(f)$ denote the unique outcome in X for which $(\mathbf{0}\,, c_s^\sigma(f)) \sim (\sigma\,, f_{\sigma^{-1}(s)} c_s^\sigma(f))$.

Since Axiom 1 implies $(\mathbf{0}\,, c_s^\sigma(f)) \sim (\sigma\,, c_s^\sigma(f))$, it follows from the transitivity of $\sim$ (which itself is implied by the existence of certainty equivalent static signal-acts that is guaranteed by Assumption 1) and the definition of $\succsim_s^\sigma$ that $f \sim_s^\sigma c_s^\sigma(f)$. That is, $c_s^\sigma(f)$ is the *conditional certainty equivalent* of the act f when the realization of the signal σ is s.

It will then be enough to show

$$v(c_s^\sigma(f)) = \int u\left(\underline{f_\sigma^\mu}\,, \overline{f}_\sigma^\mu\right)\, d\mu_s^\sigma\,,$$

where v is the Bernoulli utility obtained from u by setting $v(x) := u(x,x)$. That is, each relation $\succsim_s^\sigma$ admits an EUU representation characterized by the prior μ_s^σ and the interval utility u.

To see why this is sufficient consider the (auxiliary) act f_σ defined by setting

$$f_\sigma := \begin{cases} c_s^\sigma(f) & \text{if } \omega \in \sigma^{-1}(s) \text{ and } s \in \sigma(\Omega)^+ \\ \ell & \text{otherwise} \end{cases}$$

From the definition of the conditional certainty equivalents and by invoking Axiom 1 we have

$$(\sigma\,, f) \sim (\sigma\,, f_\sigma) \sim (\mathbf{0}\,, f_\sigma)\ .$$

And so, $V(\mathbf{0}, f_\sigma)$ represents $\succsim$.

Establishing $v \circ c_s^\sigma$ admits an EUU representation.

I deem an event E to be *ideal* with respect to the relation $\succsim_s^\sigma$ if for every pair of acts f and g:

$$f \succsim_s^\sigma g_E f \implies f_E g \succsim_s^\sigma g \quad \text{and} \quad f \succsim_s^\sigma f_E g \implies g_E f \succsim_s^\sigma g\,.$$

[13] Thus we can interpret $\succsim_s^\sigma$ as the decision-maker's interim preferences over acts when the realization of the signal σ is s.

In the terminology of Gul and Pesendorfer (2014), the first (respectively, second) implication corresponds to E being "left-ideal" (respectively, "right-ideal").[14]

Lemma 1. *Every event in the collection* $\mathcal{E}_s^\sigma = \{E \cap \sigma^{-1}(s) \colon E \in \mathcal{E}\}$ *is ideal with respect to* $\succsim_s^\sigma$*.*

Proof. Fix $E \in \mathcal{E}$. From the definition of $\succsim_s^\sigma$ and Axiom 2 we have for any pair of acts $\widehat{f}$ and $\widehat{g}$

$$\widehat{f} \succsim_s^\sigma \widehat{g} \iff \Big(\sigma\,, \widehat{f}_{\sigma^{-1}(s)} h\Big) \succsim \big(\sigma\,, \widehat{g}_{\sigma^{-1}(s)} h\big)$$

for any act h. Hence

$$\begin{aligned} f \succsim_s^\sigma g_E f &\iff \big(\sigma\,, f_{\sigma^{-1}(s)} f\big) \succsim \big(\sigma\,, (g_E f)_{\sigma^{-1}(s)} f\big) \\ \text{(applying Axiom 3)} &\implies \Big(\sigma\,, (f_E g)_{\sigma^{-1}(s)}\, f\Big) \succsim \big(\sigma\,, g_{\sigma^{-1}(s)} f\big) \\ &\iff f_E g \succsim_s^\sigma g \end{aligned}$$

That is, $E \cap \sigma^{-1}(s)$ is left-ideal. A similar construction establishes it is right-ideal as well. ■

Assumption 1 and Axioms 1 – 3 together entail the restriction of $\succsim_s^\sigma$ to acts that on the information cell $\sigma^{-1}(s)$ are adapted to $\mathcal{E}_s^\sigma$ yield a standard expected utility representation characterized by a probability μ_s^σ defined on $\mathcal{E}_s^\sigma$ and a Bernoulli utility $v_s^\sigma \colon X \to \mathbb{R}$.

Axiom 4 implies that for each non-null realization s of σ, the conditional certainty equivalent of any diffuse act on an ideal (wrt $\succsim_s^\sigma$) subset of the information cell $\sigma^{-1}(s)$ is the same as the certainty equivalent of a diffuse act for the static preferences. Hence, for each outcome interval $[x, y]$ we can set $u_s^\sigma(x, y) := u(x, y)$. And we have established that $\succsim_s^\sigma$ admits an EUU representation with the prior μ_s^σ defined on $\mathcal{E}_s^\sigma$ and interval utility $u(\cdot, \cdot)$.

It remains to show that μ_s^σ is the probability corresponding to the realization s in the uniform conditional belief system associated with the signal as specified in definition 4.

[14] It follows from Lemma B0 in Gul and Pesendorfer (2014, p25) that both implications holding is equivalent to an event satisfying their definition of an ideal event.

Since μ is convex-ranged, it readily follows for any pair of ideal events $\widehat{E}, E \in \mathcal{E}$, and any natural number $n \in \mathbb{N}$, there exists an n-element ideal partition of $\widehat{E}$, $\{\widehat{E}_i \in \mathcal{E} \colon i = 1, \ldots, n\}$ such that

$$\mu\left(\widehat{E}_i\right) = \frac{\mu\left(\widehat{E}\right)}{n} \text{ and } \mu\left(\widehat{E}_i \cap E\right) = \frac{\mu\left(\widehat{E} \cap E\right)}{n}.$$

Hence, it follows any pair of ideal events $E, E' \in \mathcal{E}$, there exists a pair of uniform refinement ideal partitions, denoted by $\{E_s \colon s \in \sigma(\Omega)^+\}$ and $\{E'_s \colon s \in \sigma(\Omega)^+\}$, respectively, that are both associated with the signal σ and the former (respectively, the latter) is associated with E (respectively, E') as specified in definition 7. Employing the expected utility representations of the static and conditional preferences over (conditionally) ideal bets with stakes m and ℓ, yields

$$V\left(\mathbf{0}, m_{E \cap E_s} \ell\right) = \mu(E \cap E_s) \text{ and } V\left(\sigma, m_{E \cap \sigma^{-1}(s)} \ell\right) = \mu_s^\sigma\left(E \cap \sigma^{-1}(s)\right)$$
$$V\left(\mathbf{0}, m_{E' \cap E'_s} \ell\right) = \mu(E' \cap E'_s) \text{ and } V\left(\sigma, m_{E' \cap \sigma^{-1}(s)} \ell\right) = \mu_s^\sigma\left(E' \cap \sigma^{-1}(s)\right).$$

Hence Axiom 5 implies :

$$\mu_s^\sigma\left(E \cap \sigma^{-1}(s)\right) \geqslant \mu_s^\sigma\left(E' \cap \sigma^{-1}(s)\right) \iff \mu(E \cap E_s) \geqslant \mu(E' \cap E'_s)$$

As this holds for all pairs of ideal events, this in turn entails that $\mu_s^\sigma(E \cap \sigma^{-1}(s))$ is proportional to $\mu(E \cap E_s)$. And since from the construction of the uniformly-distributed uniform refinement partition we have

$$\mu(E \cap E_s) = \sum_{Q \subseteq \sigma(\Omega)^+ \colon s \in Q} \frac{\mu(E \cap E_s^\sigma)}{|Q|};$$

dividing the right side of the above equation by the constant of proportionality:

$$p_s^\sigma := \mu(E_s) = \sum_{Q \subseteq \sigma(\Omega)^+ \colon s \in Q} \frac{\mu(E_s^\sigma)}{|Q|},$$

yields,

$$\mu_s^\sigma\left(E \cap \sigma^{-1}(s)\right) = \frac{\sum_{Q \subseteq \sigma(\Omega)^+ \colon s \in Q} \mu(E_s^\sigma) / |Q|}{p_s^\sigma},$$

as required. ■

The proof of the necessity of the axioms for a preference relation to admit an REUU representation is straightforward and thus omitted.

Proof of Proposition 2

Fix a non-null ideal event $G \in \mathcal{E} \backslash \mathcal{N}$. As μ is convex-valued we can construct the following sequence of ideal partitions of G, $(\mathbb{G}^1, \ldots, \mathbb{G}^n, \ldots)$.

1. $\mathbb{G}^1 := \{G_0^1, G_1^1\}$ where $\mu(G_0^1) = \mu(G_1^1) = \mu(G)/2$
2. $\mathbb{G}^2 := \{G_{00}^2, G_{00}^2, G_{10}^2, G_{11}^2, \}$ where $G_0^2 = G_{00}^2 \cup G_{01}^2$, $G_1^1 = G_{10}^2 \cup G_{11}^2$ and $\mu(B) = \mu(G)/4$ for each $B \in \mathbb{G}^2$.

   $\vdots$

n. $\mathbb{G}^n := \{\ldots, G_{b0}^n, G_{b1}^n \ldots\}$, where b is an $n-1$ bit binary number, and where for each b, $G_b^{n-1} = G_{b0}^n \cup G_{b1}^1$ and $\mu(G_{b0}^n) = \mu(G_{b1}^1) = \mu(G) \times 2^{-n}$

   $\vdots$

Gul and Pesendorfer (2014) show if the continuum hypothesis holds, then there exists a collection of pairwise disjoint diffuse events $\{D_1, D_2, \ldots\}$ such that $\bigcup_i D_i = \Omega$. Consider the following sequence of partitions of Ω.

1. $\mathbb{P}^1 := \{G_0^1 \cup \bigcup_i (D_{2i-1} \cap (\Omega \backslash G)), E_1^1 \cup \bigcup_i (D_{2i} \cap (\Omega \backslash G))\}$

2. $\mathbb{P}^2 :=$

$$\left\{ G_{00}^2 \cup \bigcup_i (D_{4i-3} \cap (\Omega \backslash G)), G_{01}^2 \cup \bigcup_i (D_{4i-2} \cap (\Omega \backslash G)), \right.$$

$$\left. G_{10}^2 \cup (D_{4i-1} \cap (\Omega \backslash G)), G_{11}^2 \cup \bigcup_i (D_{4i} \cap (\Omega \backslash G)) \right\}$$

   $\vdots$

Notice each partition in this sequence is a refinement of its predecessor and for each n, each element in the partition $\mathbb{P}^n$ has inner measure $\mu(G)\gamma(2^{-n})$. So consider the sequence of σ-algebras $(\widehat{\mathcal{E}}^n)_{n=1}^{\infty}$ in which $\widehat{\mathcal{E}}^n$ is the σ- algebra generated by the partition $\mathbb{P}^n$. For each n and each $A \in \widehat{\mathcal{E}}^n$, set

$$\pi^n(A) := \frac{\mu(A)}{\mu(G)}.$$

By construction, we have for each $A \in \widehat{\mathcal{E}}^n$,

$$\mu_*(A) = \begin{cases} \mu(G)\,\pi^n(A) & \text{if } \pi^n(A) < 1 \\ 1 & \text{otherwise} \end{cases}$$

Finally, take $(\pi, \mathcal{E}^\pi)$ to be the completion of $(\pi^\infty, \bigcup_n \widehat{\mathcal{E}}^n)$, where π^∞ is the measure on $\bigcup_n \widehat{\mathcal{E}}^n$ obtained by setting $\pi^\infty(A) := \pi(A)$ for every $A \in \bigcup_n \widehat{\mathcal{E}}^n$.
■

References

Cerreia-Vioglio, S., Corrao, R., Lanzani, G., January 2024. (un-)common prreferences, ambiguity, and coordination.

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

Dempster, A. P., 1967. Upper and Lower Probabilities Induced by a Multivalued Mapping. Annals of Mathematical Statistics 38, 325–339.

Denneberg, D., 1994. Conditional (updating) non-additive measures. Annals of Operation Research 52, 21 – 42.

Denneberg, D., 2002. Conditional expectation for monotone measures, the discrete case. Journal of Mathematical Economics 37, 105 – 121.

Eichberger, J., Kelsey, D., Grant, S., 2007. Updating choquet beliefs. Journal of Mathematical Economics 43, 888–899.

Epstein, L. G., Schneider, M., 2003. Recursive multiple-priors. Journal of Economic Theory 113, 1–31.

Epstein, L. G., Zin, S. E., 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57 (4), 937–969.

Frisch, D., Baron, J., 1988. Ambiguity and Rationality. Journal of Behavioral Decision Making 1, 149–157.

Ghirardato, P., 2002. Revisiting Savage in a conditional world. Economic Theory 20, 83–92.

Gilboa, I., Schmeidler, D., 1993. Updating ambiguous beliefs. Journal of Economic Theory 59, 33–49.

Grant, S., Kajii, A., Polak, B., 1998. Intrinsic preference for information. Journal of Economic Literature 83, 233–259.

Grant, S., Liu, S. L., Yang, J., January 2024. A Comment on "Expected Uncertain Utility". Econometrica 92 (1), 247–256.

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

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

Gul, F., Pesendorfer, W., 2021. Evaluating ambiguous random variables from Choquet to maxmin expected utility. Journal of Economic Theory 192, 1–27.

Hanany, E., Klibanoff, P., 2007. Updating preferences with multiple priors. Theoretical Economics 2 (3), 261–298.

Jaffray, J.-Y., 1992. Bayesian updating and belief functions. IEEE Transactions on System, Man and Cybernetics 22, 1144–1152.

Kreps, D. M., Porteus, E. L., 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. Econometrica 46 (1), 185–200.

Lapied, A., Kast, R., Toquebeuf, P., 2012. Updating Choquet capacities: a general framework. Economics Bulletin 2012 32 (2), 1495 – 1503.

Lehrer, E., 2005. Updating non-additive probabilities – a geometric approach. Games and Economic Behavior 50, 42 – 57.

Machina, M. J., 2009. Risk, ambiguity, and the rank-dependence axioms. American Economic Review 99 (1), 385–392.

Machina, M. J., Schmeidler, D., 1992. A More Robust Definition of Subjective Probability. Econometrica 60 (4), 745–780.

Piermont, E., 2021. Unforeseen evidence. Journal of Economic Theory 193, 105235.

Pires, C., 2002. A rule for updating ambiguous beliefs. Theory and Decision 53, 137–152.

Sarin, R. K., Wakker, P. P., 1998. Dynamic choice and non-expected utility. Journal of Risk and Uncertainty 17, 87–119.

Shafer, G., 1976. A Mathematical Theory of Evidence. Princeton University Press, Princeton, New Jersey.

Siniscalchi, M., 2011. Dynamic choice under ambiguity. Theoretical Economics 6, 379–421.

Walley, P., 1991. Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London.

Wang, T., 2003. Conditional preferences and updating. Journal of Economic Theory 108, 286–321.