Decision rules for imprecise Lockeans

Hana Kalpak

intro

This talk is about how our **coarse-grained** doxastic attitudes relate to our **fine-grained** doxastic attitudes, if we are rational.

By **coarse-grained doxastic attitudes**, I mean the triple of

- **belief**: the attitude we have towards things we *take to be the case*,
- **disbelief**: the attitude we have towards things we *take not to be the case*,
- **suspension of judgment**: the attitude we have towards things we *neither take to be the case, nor not to be the case*.

By **fine-grained doxastic attitudes**, I mean **credences**: the various *degrees of confidence* we can have in somethings' being the case, ranging from complete confidence to complete lack of confidence.

Two assumptions to get us off the ground:

1. Whenever we have a rational doxastic attitude of one granularity towards something, there is some rational doxastic attitude of the other granularity towards that something;
2. If we are rational, there are certain limitations on which pairs of coarse- and fine-grained doxastic attitudes we can have towards the same thing.

Intuitively: If you have very low confidence in something's being the case, it seems **irrational** for you to believe it.

We will be concerned with a question raised by assumption #2:

What is the general property distinguishing rational combinations of coarse- and fine-grained doxastic attitudes from irrational ones?

A version of a much-discussed type of response (Foley, 1993):

The Lockean thesis (belief version). There is some non-extremal credence r, such that for any proposition p: one's belief in p is rational *if and only if* one's (rational) credence in p exceeds r.

That is: Rational belief corresponds exactly to rational credence above some threshold value between the extremes (note the wide-scope existential).

Pairing this with a similar threshold for **suspension of judgment** (below which one **disbelieves**), we get a general story about the relation between rational coarse-grained and rational fine-grained doxastic attitudes.

intro

The Lockean thesis has some well-known drawbacks:

- Given the assumption that rational credences are **probabilistic**, the thesis violates principles of *closure* and *completeness* of beliefs;
- If supplemented to adjust for this, we allow for *lottery*- and *preface paradoxes*.

I will present a separate challenge to the Lockean thesis, resulting from an alternative view of what our fine-grained doxastic attitudes look like.

intro

Standardly, our fine-grained attitudes are treated as infinitesimally **precise**.

Your credence in a proposition p is represented as a point-value from $[0, 1]$, given by a credence function c:

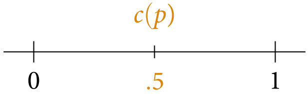

According to another tradition, our fine-grained attitudes may be **imprecise**.

An **imprecise credence** in a proposition p can be represented by a subinterval of $[0, 1]$, given by a set C of credence functions:

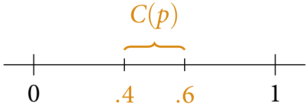

If rational credences can be imprecise, the Lockean thesis breaks down.

The problem: We lack an intuitive connected order on intervals: $[.4, .6]$ is neither clearly greater than, lesser than, or equal to things like $[.5, .5] = .5$.

Assume the belief threshold is .5, and you have a rational $[.4, .6]$ credence in p. Then the Lockean thesis deems any coarse-grained doxastic attitude towards p irrational, contradicting our assumption #1.

The question for the Lockean:

When our rational credence is imprecise, how are our rational coarse-grained doxastic attitudes determined?

intro

The discussion will be framed within a particular **utility-theoretic** understanding of the Lockean thesis.

By assigning utilities to coarse-grained doxastic attitudes reflecting how **accurately** they represent the world, Lockean thresholds can be derived from the assumption that these attitudes are rational *iff* they *maximize expected utility* given one's rational credences.

On this formulation of the Lockean thesis, the problem of accommodating imprecision becomes a problem of finding a suitable **decision rule** to ground the relation between rational fine- and coarse-grained doxastic attitudes.

intro

What we will do today:

- Extend a version of **epistemic utility theory** to include imprecise fine-grained doxastic attitudes alongside coarse-grained ones,
- formulate some desiderata for rational decision rules within imprecise epistemic utility theory, and
- look at how a couple of rules from practical imprecise decision theory fare with respect to these desiderata.

What we won't do today: Find a rule that satisfies all these desiderata. This is work in progress! The main contribution of the talk is instead to set the stage for further investigation both of suitable decision rules and suitable assessment criteria.

roadmap

- Utility-theoretic Lockeanism
  - *Epistemic utility theory, utility-theoretic Lockeanism*
- Imprecise credences
  - *Motivation, imprecise epistemic utility theory.*
- Desiderata for imprecise epistemic decision rules
  - *Basic properties, unanimity, safety, persistence*
- Rules for imprecise epistemic choice
  - *E-admissibility,* Γ*-maximin, composite rules*
- Conclusion
- References
- Proofs

Utility-theoretic Lockeanism

epistemic utility theory

Epistemic utility theory (EUT) uses methods from decision theory to study rationality criteria for doxastic attitudes.

The term is typically associated with works aiming to establish Bayesian rationality criteria for credences.[1]

An EUT-approach to combinations of coarse- and fine-grained doxastic attitudes was championed by Carl Hempel (1962), and revived in Easwaran (2016) and Dorst (2019)'s derivations of a Lockean thesis from a certain version EUT.

To see how this works, we begin by outlining a basic EUT-framework recognizing both coarse- and fine-grained doxastic attitudes.

[1] For instance, Joyce, 1998; 2009; Greaves & Wallace, 2006; Leitgeb & Pettigrew, 2010ab; Pettigrew, 2016.

epistemic utility theory

EUT is motivated by a general consequentialism about epistemic rationality: a doxastic attitude is rational insofar as adopting it is a rational means of maximizing the **epistemic value** of one's doxastic state.

Different specifications of "epistemic value" yield different versions of EUT.

According to an especially prominent version, which we will follow here, **epistemic value = accuracy**: our doxastic attitudes are *good* to the extent that they *accurately reflect the world*.

Thus, the value of a belief or high credence in a truth is higher than the value of a disbelief or low credence in it, and vice versa for a falsehood.

Given a specification of epistemic value, our doxastic states can be assigned **utilities** (numerical values), reflecting *the epistemic value of being in that state, given what the world is like*.

Doxastic states can then be treated as decision-theoretic **options** which may be more or less rational for an agent to adopt, given *the utilities of her options*, her *expectations about the world*, and her *appetite for risk*.

Being in a particular doxastic state can then be seen as having adopted that state over others, according to a more or less rational **decision rule**.

The rationality of a belief state, then, is ultimately determined by **the rationality of the decision rule** one would use to pick that state over others.[2]

[2] Strictly, "would if possible": we do not assume *voluntarism*, i.e. that we actually *can* choose our beliefs, in such a direct manner.

EUT gives us a framework for theorizing about rational choice in the context of **epistemic decision problems** (e.d.p:s).

Like a practical decision problem (under uncertainty), an e.d.p. involving a single agent a is characterized by

- a space O of **options** among which a is to choose,
- a space W of **ways** in which the world can be, and
- an agent-specific **utility function** $U_a : O \times W \mapsto \mathbb{R}$.

They are distinguished from corresponding practical decision problems by certain requirements put on the **options** and on the **utility function**.

In an e.d.p., an **option space** is the power set of an **agenda** $\mathcal{A}$: a set of propositions[3] (closed under negation, for simplicity), intuitively corresponding to a set of propositions that are *under consideration* for some agent, at a time.

An **option** is a **belief set**: a set $\mathcal{B}$ of propositions, intuitively corresponding to a set of propositions that an agent has the option of jointly believing, given the ones she has under consideration.

Belief sets represent the coarse-grained doxastic attitudes of an agent. Where $\mathcal{B}$ is her belief set and p a proposition on her agenda:

- If $p \in \mathcal{B}$, we say that the agent **believes** p,
- if $\neg p \in \mathcal{B}$, we say that the agent **disbelieves** p, and
- if $p, \neg p \notin \mathcal{B}$, we say that the agent **suspends judgment** on p.

[3] Propositions are given the standard treatment as *sets of possible worlds*; true at the worlds they contain, and false at the others.

The fine-grained doxastic attitudes of an agent a are represented by a **credence function** $c_a : \mathcal{A} \mapsto [0,1]$, where the size of the value c_a assigns to a proposition represent the strength of a's confidence in the proposition.

Ultimately, we will be interested in how rational fine-grained attitudes constrain rational coarse-grained ones. Thus, we can presuppose that credences are rational, at least in the sense of being **probabilistic**.[4]

Probabilism. Credence functions are probability functions: for any c_a,

- $c_a(W) = 1$ and $c(\emptyset) = 0$,
- $c_a(p \cup q) = c_a(p) + c_a(q)$, for any disjoint p, q in the domain of c_a.

[4] We may assume that the alternatives have already been ruled out by an accuracy-dominance argument for probabilism in the style of Joyce (1998).

The **utility function** of an e.d.p. is defined using an **accuracy measure**.

An **accuracy measure** for coarse-grained doxastic states is a function that assigns a **true belief** some value V^T, a **lack of belief** some strictly lesser value V^S, and a **false belief** some even lesser value V^T (all in $\mathbb{R}$).

That is: At a way w, the accuracy of a doxastic attitude towards a proposition p (encoded by a belief set $\mathcal{B}$) is given by

$$v(\mathcal{B}, p, w) = \begin{cases} V^T & \text{if } p \in \mathcal{B} \text{ and } p \text{ is } \textit{true} \text{ at } w \\ V^S & \text{if } p, \neg p \notin \mathcal{B} \\ V^F & \text{if } p \in \mathcal{B} \text{ and } p \text{ is } \textit{false} \text{ at } w \end{cases}$$

where $V^T > V^S > V^F$.

We use accuracy measures to define the epistemic utility of a belief set $\mathcal{B}$ as **the sum of the accuracy of the individual coarse-grained attitudes** encoded by $\mathcal{B}$, given a way w in which the world might be.

Formally, this value is given by an agent's (a) epistemic utility function U_a:

$$U_a(w, \mathcal{B}) := \sum_{p \in \mathcal{A}} v(\mathcal{B}, p, w)$$

Choices between doxastic states are choices under **uncertainty**: we typically lack knowledge about the *objective chance* of our beliefs' being accurate.

Still, our credences encode **expectations** about the accuracy of our beliefs.

An **expected epistemic utility function** E_{c_a} outputs values encoding the **expected accuracy** of a belief state, given the agent's expectations about the world (c_a):

$$E_{c_a}(\mathcal{B}) = \sum_{w \in W} c_a(\{w\}) U_a(w, \mathcal{B})$$

That is: we get the expected epistemic utility (accuracy) of a belief state by calculating, for each way w, the utility of the belief state at w, weigh this by the credence assigned to $\{w\}$, and sum the results.

A constraint on rational coarse-grained attitudes can now be captured as a version of the *expected utility hypothesis*:

Maximize Expected Accuracy. A belief set is rational iff it maximizes expected accuracy (epistemic utility).

From this assumption, we will be able to derive a Lockean thesis.

A point on notation. An e.d.p. is characterized by a triple $(2^{\mathcal{A}}, W, U_a)$. Assuming W is fixed, we can write this more economically: from now on, $\mathcal{A}_a$ is short for $(2^{\mathcal{A}}, W, U_a)$.

Adam (a) glimpses a shadow in the water, and for the first time comes to consider whether p = *There are sharks in the Baltic Sea*.

Fun facts about Adam:

- his credences are probabilistic,
- his belief set maximizes expected accuracy,
- he values a true belief just as much as he disvalues false one, relative to the value of suspending judgment.

Without loss of generality, we can define Adam's epistemic utility function U_a using an accuracy measure $v : V^T = 1, V^S = 0, V^F = -1$.

The problem $\{p, \neg p\}_a$ can be represented in a **decision table**:

		$\mathcal{B}_1$ $\{p\}$	$\mathcal{B}_2$ $\{\neg p\}$	$\mathcal{B}_3$ $\{\}$
w_1	p	1	-1	0
w_2	$\neg p$	-1	1	0

Reading guide: $\mathcal{B}_1 = \{p\}, \mathcal{B}_2 = \{\neg p\}$ and $\mathcal{B}_3 = \{\}$ are the three belief sets constituting Adam's options, and w_1, w_2 are ways the world would be if ***there are sharks in the Baltic Sea*** (w_1) and if ***there are no sharks in the Baltic Sea*** (w_2). Any other way is analogous to either w_1 or w_2, so can be ignored.

For $i, j \in \{1, 2, 3\}$: the value at $(w_i, \mathcal{B}_j)$ is the **epistemic utility** Adam receives from the belief set $\mathcal{B}_j$ at w_i (either V^T, V^S or V^T since $\{p, \neg p\}_a$ is "atomic").

epistemic utility theory · illustration

This graph shows how Adam's **epistemic utility** (y-axis) for the three options in $\{p, \neg p\}_a$ varies with the **credence** he assigns to p (x-axis).

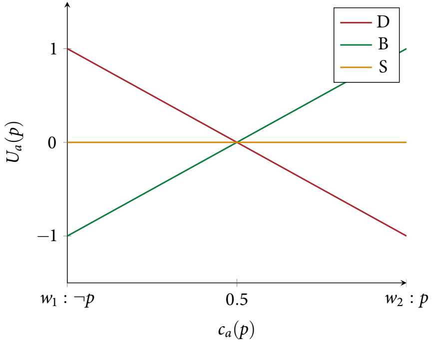

If $c_a(p) < .5$, Adam will **disbelieve** p (D). If $c_a(p) > .5$, he will **believe** p (B).
If $c_a(p) = .5$, he might do either, or **suspend judgment** on p (S).

In this simple example, **the relation between Adam's coarse- and fine-grained doxastic attitudes conform to the Lockean thesis.**

There is some non-extremal credence r (namely, .5), such that Adam's belief in a proposition is rational *if and only if* his rational credence is at or above r.

Generally: $\mathcal{B}$ will be rational for an agent a with agenda $\mathcal{A}$ *iff*, for all $p \in \mathcal{A}$:

- if $p \in \mathcal{B}$, then $c_a(p) \geq \frac{-V^F}{V^T - V^F}$, bt
- if $\neg p \in \mathcal{B}$, then $c_a(p) \leq 1 - \frac{-V^F}{V^T - V^F}$, and st
- if $p, \neg p \notin \mathcal{B}$, then $(1 - \frac{-V^F}{V^T - V^F}) \geq c_a(p) \leq \frac{-V^F}{V^T - V^F}$.

Thus by assuming *Maximize Expected Accuracy*, we get:

Tripartite Lockean thesis. For any proposition p on your agenda:

- **belief** in p is rational *if and only if* your credence in p is at least bt,
- **disbelief** in p is rational *if and only if* your credence in p is at most st,
- **suspending judgment** in p is rational *if and only if* your credence in p is at least st and at most bt.

In Adam's case, bt and dt coincide: his utility for having a true belief **equals** his disutility for having a false belief. This in effect models him as **risk neutral**.

We can model agents with different appetites for doxastic risk by specifying the **ratio** between the outputs of the accuracy measure:

- U_a is NEUTRAL iff $|V^T| - |V^F| = |V^S|$.
  - Example: $V^T = 1$, $V^F = -1$, $V^S = 0$.
- U_a is CONSERVATIVE iff $|V^T| - |V^F| < |V^S|$.
  - Example: $V^T = 1$, $V^F = -1.5$, $V^S = 0$.
- U_a is RADICAL iff $|V^T| - |V^F| > |V^S|$.
  - Example: $V^T = 1.5$, $V^F = -1$, $V^S = 0$.

Not all of these profiles are intuitively rational in themselves (especially, RADICAL agents may prefer believing contradictory propositions), but that's a separate issue. Still, I will focus on NEUTRAL and CONSERVATIVE agents.

To sum up, we have seen that the following assumption yields a tripartite Lockean thesis, generalized over risk profiles:

Maximize expected accuracy. A belief set is rational for an agent *a* iff it maximizes expected accuracy (epistemic utility) according to *a*'s rational credences.

Imprecise credences

What we have said so far will help you assess the rationality of your coarse-grained doxastic attitudes in many everyday scenarios.

Fair coin. The Oracle shows you a coin, tells you that the coin is fair, and tosses it. You rationally lend a .5 credence to the event that ***the coin lands heads-up*** (HEADS). Given this, should you believe HEADS, disbelieve HEADS, or suspend judgment on the matter?

If rational doxastic attitudes maximize expected accuracy, you are rational in believing HEADS *iff* your $\mathsf{bt} \geq .5$, rational in disbelieving *iff* your $\mathsf{st} \leq .5$, and rational in suspending *iff* $\mathsf{bt} \leq .5 \geq \mathsf{st}$.

But what about these types of scenarios?

> **Mystery coin.** The Oracle shows you a new coin, and tells you only that the coin is fair *or* biased to some degree, in some direction. She tosses it, while you consider the proposition that ***the mystery coin lands heads-up*** (HEADS*). Should you believe HEADS*, disbelieve it, or suspend judgment?

Like before, this depends on the credence you lend HEADS*.

Unlike before, the rational credence to lend cannot be determined by appeal to your beliefs about objective chance.

For all you know, the objective chance of HEADS* may lie *anywhere in between* 0% and 100%. You have no evidence favoring any one option above another.

So: **What would be a rational credence in HEADS*?**

According to the classic Bayesian, the rational credence is a member of [0, 1], weak evidence notwithstanding.

- **Subjectivist reasoning:** If the evidence is balanced over a set of credences, each is equally good, and you may rationally pick whichever is to your subjective liking.
- **Objectivist reasoning:** By the *principle of indifference*, you must have a .5 credence in HEADS*, since you have no more evidence of HEADS* than you have of ¬HEADS*.

Still, assigning a credence $x \in [0, 1]$ to HEADS* will not capture the **severity** of your uncertainty about the proposition.

Illustration. Say that $c_y(\textsc{heads}*) = .5$ (but note that the argument can be made perfectly general). Coincidentally, the exact same credence as you had in HEADS.

If a rational agent (like you) assign two propositions the **same credence**, she would standardly be expected to be **equally willing to bet** on their respective truth.

But should you have the option to place a bet on the same amount for either HEADS or HEADS*, it is not obviously irrational to have a preference for one over the other.

In particular, many have the intuition that a *preference for the fair coin* is rational (related: Ellsberg, 1961).

The classic Bayesian response cannot explain such a preference.

An alternative line of response is to deny that your rational credence in HEADS* can be identified with any particular value from $[0,1]$ (e.g., Joyce, 2010).

Instead, we identify this credence with a **set** of values from $[0,1]$: the set of values corresponding to all objective chances of HEADS* you lack reason to exclude.

Modeling rational credences in this way distinguishes your credence in HEADS ($= .5$, or $\{.5\}$) from your credence in HEADS* ($= [0,1]$), on the basis of the noted differences of your evidence between the two cases.

The idea that rational credences can be *imprecise*—identified with sets of values, rather than with single values—is explored and defended in a wide range of works from the last century, including Keynes (1921), Gärdenfors & Sahlin (1982), Levi (1980, etc.), Walley (1992), Joyce (2010).

Imprecise credences might offer a way to model how our fine-grained doxastic attitudes reflect the **weight** and **balance** of our evidence (without invoking things like higher-order credences).

Other motivations come from work on the **incommensurability of options** (Keynes, Levi), **ambiguity aversion** (Ellsberg), **unknown correlations** (Haenni et al., 2011) and **suspension of graded judgment** (Walley, 1991).

If rational credences can be imprecise in this sense, the EUT we have outlined so far cannot provide a complete story about the relation between rational fine- and coarse-grained attitudes.

Maximize Expected Accuracy is defined for **precise expected utilities**, computed from precise credences.

The simplest way of calculating expected utilities from imprecise credences will instead yield **imprecise expected utilities**: intervals, which are not generally maximizable.

In other words, the EUT we have considered handles only **precise e.d.p.:s**: ones where **all options all have precise expectation values**.

To handle cases like **Mystery coin**, we need an EUT for **imprecise e.d.p.:s**: ones where **some options have imprecise expectation values**.

Following Isaac Levi (a long-time proponent of imprecise credences), we will model imprecise credences as values given by **credal sets** (Levi, 1974 etc.).

A **credal set**, denoted C_a for an agent a, is a **closed and convex** set of credence functions, meaning that it

- has **extreme points**, and *closure*
- contain **all linear averages** of their members. *convexity*

The **imprecise credence** assigned to a proposition p is simply the set

$$C_a(p) = \{c(p) \mid c \in C_a\}.$$

Given our assumptions, any $C_a(p)$ is a closed subinterval of $[0, 1]$.

By taking all the regular expected utilities calculated from the members of some credal set, we get an **imprecise expected epistemic utility**.

That is: Define the imprecise expected utility of $\mathcal{B}$ given C_a, denoted by $IE_{C_a}(\mathcal{B})$, as

$$IE_{C_a}(\mathcal{B}) := \{E_c(\mathcal{B}) \mid c \in C_a\}.$$

Given our definition of credal sets, imprecise expected utilities are closed, convex subintervals of $\mathbb{R}$.

I will use $\underline{IE}_{C_a}(\mathcal{B})$ and $\overline{IE}_{C_a}(\mathcal{B})$ to denote the left and right endpoint, respectively, of an imprecise expected utility $IE_{C_a}(\mathcal{B})$.

We now have the general components of an **imprecise epistemic utility theory**: a definition of imprecise e.d.p.:s and a definition of imprecise expected utility.

To assess the rationality of an option in an imprecise e.d.p., we still need a **decision rule** defined for imprecise expected utilities.

Before looking at some rules of this kind, I will suggest some **desiderata**: properties that seem good for such rules to have, given that they are to be used for specifically **epistemic** decision problems (precise or imprecise).

Desiderata for imprecise epistemic decision rules

desiderata · choice functions

To formulate and investigate the formal properties of our candidate decision rules, it is useful to formulate them in terms of choice functions.[5]

A **choice function** f takes a decision problem $\mathcal{A}_a$ and outputs the subset of $2^{\mathcal{A}}$ whose members satisfy some certain criteria.

We say that

- $f(\mathcal{A}_a)$ is the set of permitted options in $\mathcal{A}_a$, according to f, and
- $2^{\mathcal{A}} - f(\mathcal{A}_a)$ is the set of rejected options in $\mathcal{A}_a$, according to f.

When a choice function f permits all and only the options that conform to to some decision rule D, I will say that f characterizes D.

[5] I use this term somewhat idiosyncratically, but innocently so.

Maximize Expected Accuracy. A belief set is rational iff it maximizes expected accuracy (epistemic utility).

Maximize Expected Accuracy is characterized by the choice function MAX:

$$\textsc{max}(!\mathcal{A}_a) = \{B \subseteq \mathcal{A} \mid \forall c \in C_a : \forall \mathcal{B}' \subseteq \mathcal{A} : E_c(B) \geq E_c(\mathcal{B}')\}$$

MAX takes a precise decision problem $!\mathcal{A}_a$ and outputs the set of belief sets in $2^{\mathcal{A}_a}$ that maximize expected accuracy according to the members of C_a.[6]

[6] In a precise problem these all agree, so can be represented by any single member.

desiderata · basic properties

All rules to be considered are characterized by choice functions f that are

- **Decisive.** For any $\mathcal{A}_a$: $f(\mathcal{A}_a) \neq \emptyset$.
  There is always some rational belief set.
- **Fixating.** For any $\mathcal{A}_a$: If $f(\mathcal{A}_a)$ is an agenda, then $f(f(\mathcal{A}_a)_a) = f(\mathcal{A}_a)$.
  All rational belief sets are equally rational.
- **Coherent.** For any $(!\mathcal{A}_a)$: $f(!\mathcal{A}_a) = \textsc{max}(!\mathcal{A}_a)$.
  When expectation values are precise, a belief set is rational iff it maximizes expected accuracy.

I take these to be uncontroversially desirable properties of imprecise epistemic decision rules.

All of the decision rules to be considered are also **unanimous**: characterized by choice functions with the following property.

Unanimity. If $E_c(\mathcal{B}) \geq E_c(\mathcal{B}')$ for all $c \in C_a$ and $\mathcal{B}' \subseteq \mathcal{A}$, then $\mathcal{B} \in f(\mathcal{A}_a)$.

That is: if an option maximizes expected accuracy according to all members of the credal set, then it is permitted.

This guarantees that a rational belief set remains rational if one's credence is **precisified** by update of C_a, for the special case where *this specific belief set* is unanimously prefereed within C_a.

desiderata · strong unanimity

Not all rules will be **strongly unanimous**:

> **Strong unanimity.** If $[E_c(\mathcal{B}) \geq E_c(\mathcal{B}')$ for all $c \in C_a$ and $\mathcal{B}' \subseteq \mathcal{A}]$, then $f(\mathcal{A}_a) = \{\mathcal{B} \mid \forall c \in C_a : \forall \mathcal{B}' \subseteq \mathcal{A} : E_c(\mathcal{B}) \geq E_c(\mathcal{B}')\}$.

That is: if some options maximize expected accuracy according to all members of the credal set, **only these options** are permitted.

This guarantees that a rational belief set remains rational if one's credence C_a is precisified, for the sligthly more general case where *some belief set* is unanimously preferred within C_a.[7]

[7] The most general version of this desideratum, of course, does not condition on unanimity. We will return to this.

desiderata · dominance

Rational decision rules are often required to reject **dominated** options.

Dominance. An option O' is **dominated** by O whenever

- for any possible state of the world, O yields at least as much utility as O' does, and
- for some possible states of the world, O yields strictly more.

Choosing a dominated option is intuitively irrational: switching to the dominating option can only make you better off.

Fact 1. MAX rejects dominated options (assuming state-act independence).

Consequence: *All coherent rules reject dominated options in precise problems.*

desiderata · dominance

For imprecise problems, we can look at the relation of collective dominance:

Collective dominance. In a decision problem $\mathcal{A}_a$, an option O *collectively dominates* an option O' iff

- $E_c(O) \geq E_c(O')$ for all $c \in C_a$, and
- $E_c(O) > E_c(O')$ for some $c \in C_a$.

Some of the decision rules to be considered are characterized by safe choice functions:[8]

Safety. No $\mathcal{B} \in f(\mathcal{A}_a)$ is collectively dominated.

[8] Safe choice functions may still permit options dominated by *mixed* options. I'm not sure how to make sense of mixed strategies in the epistemic setting.

desiderata · persistence

Rational coarse-grained attitudes are intuitively **persistent** in a certain sense.

If you rationally believe that p, it seems rational to persist in this belief, and irrational to give it up, *as long as your credence in p remains unchanged*.

In particular, you should be able to **consider new propositions**, or **cease to consider ones**, without this itself mandating a change in attitude towards p.

Note. The motivation for this depends on how we conceptualize the agenda itself: ultimately, on what sort of doxastic attitudes we take ourselves to be modeling (occurrent? dispositional?). But least for the occurrent case, both requirements appear intuitive.

A decision rule predicts this type of persistence *iff* it is characterized by a choice function that is **e(xpansion)-consistent** and **c(ontraction)-consistent**. Where Γ is some set of propositions closed under negation:

- **E-consistency.** $f(\mathcal{A}_a \cup \Gamma) = \{\mathcal{B} \cup \mathcal{B}' \mid \mathcal{B} \in f(\mathcal{A}_a) \text{ and } \mathcal{B}' \in f(\Gamma_a)\}$.
- **C-consistency.** $f(\mathcal{A}_a - \Gamma) = \{\mathcal{B} - \Gamma \mid \mathcal{B} \in f(\mathcal{A}_a)\}$.

E-consistency ensures that expansions preserve the rational coarse-grained attitudes towards propositions from the original agenda.

C-consistency ensures that contractions preserve any rational coarse-grained attitudes towards propositions remaining from the original agenda.

A decision rule is jointly e- and c-consistent *iff* it is characterized by a **persistent** choice function:

Persistence. $\mathcal{B} \in f(\mathcal{A}_a)$ iff, for all agendas $A \subseteq \mathcal{A}$: $\mathcal{B} \cap A \in f(A_a)$.

A persistent choice function permits a belief set in an e.d.p. *iff* it permits its solutions to all subproblems of that e.d.p..

Terminology: the **$\mathcal{B}$-solution** to $\mathcal{A}_a$ is the set $\mathcal{B} \cap \mathcal{A}$. A **subproblem** of $\mathcal{A}_a$ is any problem A_a : A is an agenda included in $\mathcal{A}$.

Note that *Maximize Expected Accuracy* is persistent: $\mathcal{B} \in \textsc{max}(\mathcal{A}_a)$ *iff* $\mathcal{B} \cap A \in \textsc{max}(A_a)$ for all agendas $A \subseteq \mathcal{A}$.[9]

[9] Proof in Appendix.

desiderata · summary

There are many intuitively desirable properties of epistemic decision rules (see Conclusion). But time forces us to limit attention, and we will focus on the set of desiderata discussed so far:

- **Unanimity.** If a belief set maximizes expected accuracy according to all members of the credal set, then it is permitted.
- **Strong unanimity.** If some belief sets maximize expected accuracy according to all members of the credal set, then only these belief sets are permitted.
- **Safety.** Collectively dominated options are rejected.
- **Persistence.** Expanding or contracting an agenda does not affect the rationality of individual coarse-grained attitudes towards propositions preserved from the previous agenda.

These are thus to be seen as potentially *necessary*, but definitely *not sufficient*, properties of rational decision rules for imprecise epistemic choice.

Rules for imprecise epistemic choice

rules

We will consider two main decision rules for imprecise choice:

- *E-admissibility* (Levi 1980),
- Γ-*maximin* (Gilboa & Schmeidler 1989),

together with some composite rules based on these.

Our goal is to assess their suitability for imprecise epistemic choice, by checking whether they are

- **unanimous**,
- **strongly unanimous**,
- **safe**,
- **persistent**.

rules · e-admissibility

Isaac Levi (1974, onwards) advocates *E-admissibility* as a necessary (but perhaps not sufficient) decision rule for imprecise choice:

E-admissibility. Permitted options maximize expected utility according to some member of your credal set.

E-admissibility is characterized by the choice function E-MAX:

$$\textsc{e-max}(\mathcal{A}_a) = \{\mathcal{B} \subseteq \mathcal{A} \mid \exists c \in C_a : \forall \mathcal{B}' \subseteq \mathcal{A} : E_c(\mathcal{B}) \geq E_c(\mathcal{B}')\}$$

rules · e-admissibility

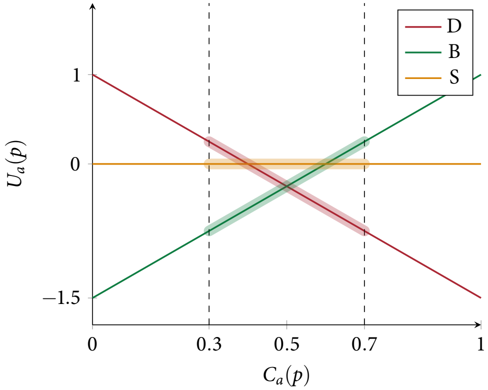

Illustration. A CONSERVATIVE risk profile with [.3, .7] credence in p. The thicker, shaded lines mark the permitted solutions according to *E-admissibility*, including **disbelief**, **suspension**, and **belief**.

Good news: *E-admissibility* is **unanimous**. If all members of your credal set agree that an option maximizes expected accuracy, this option is permitted.

Bad news: *E-admissibility* is **not strongly unanimous**. Unless all members of your credal set agree *exactly* on which options maximize expected accuracy, options that are not unanimously preferred may also be permitted.

E-admissibility is **not safe**. If there is disagreement among the members of your credal set, dominated choices may be permitted.

rules · e-admissibility

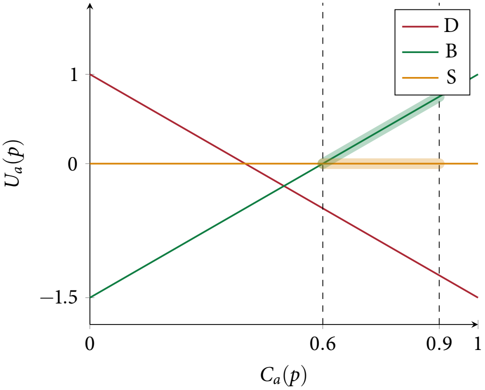

Counterexample. A CONSERVATIVE risk profile with $[.6, .9]$ credence in p. The thicker, shaded lines mark the permitted solutions according to *E-admissibility*, including **belief** but also **suspension**, which is dominated and not unanimously preferred.

More bad news: *E-admissibility* is **not persistent**.

If E-MAX permits a belief set for an e.d.p., it also permits the set's solutions to the subproblems of the e.d.p.:[10]

$$\mathcal{B} \in \text{E-MAX}(\mathcal{A}_a) \Rightarrow \mathcal{B} \cap A \in \text{E-MAX}(A_a) \text{ for all agendas } A \subseteq \mathcal{A}$$

But **the converse does not hold**: A belief set may solve all subproblems of some e.d.p., but not be permitted by E-MAX in the full e.d.p.:

$$\mathcal{B} \in \text{E-MAX}(\mathcal{A}_a) \not\Leftarrow \mathcal{B} \cap A \in \text{E-MAX}(A_a) \text{ for all agendas } A \subseteq \mathcal{A}$$

[10]Proof in Appendix.

Counterexample. Consider an e.d.p. $\mathcal{A}_a = \{p, \neg p, q, \neg q\}$, where

- C_a contains only two (probabilistic) credence functions $c_1, c_2 : c_1(p) = c_1(\neg q) = c_2(q) = c_2(\neg p) > .5$ and their mixtures;
- U_a is CONSERVATIVE, with $V^S = 0$ and $V^F = -1.5V^T$.

We then have that $\text{E-MAX}(\{p, \neg p\}_a) = \{\{p\}, \{\neg p\}, \{\}\}$: each option maximizes for the linear average of c_1 and c_2, which assigns all propositions a .5 credence. For analogous reasons, $\text{E-MAX}(\{q, \neg q\}_a) = \{\{q\}, \{\neg q\}, \{\}\}$.

Then $\mathcal{B} = \{p, q\}$ solves both subproblems of $\mathcal{A}_a$, by $\{p\}$ and $\{q\}$. Yet $\mathcal{B} \notin \text{E-MAX}(\mathcal{A}_a)$: for all $c \in C_a$, $E_c(\mathcal{B})$ will be a negative value, while $E_c(\{\})$ is always 0.

In other words, *E-admissibility* is not e-consistent: by expanding the agenda, rational agents (even when CONSERVATIVE) can be forced to change their attitude towards propositions from the original agenda.

In sum:

	e-admissibility
unanimity	✓
strong unanimity	✗
safety	✗
e-consistency	✗
c-consistency	✓

Before discussing the possibility of supplementing or strengthening *E-admissibility*, we investigate another candidate rule.

rules · Γ-maximin

Γ-maximin (e.g., Gilboa & Schmeidler, 1989) is a simple maximin-rule for imprecise decision making:

Γ-*maximin*. Permitted options maximize the lowest expected utility assigned by any member of your credal set.

Γ-*maximin* is characterized by the choice function MAXIMIN:

$$\textsc{maximin}(\mathcal{A}_a) = \{\mathcal{B} \subseteq \mathcal{A} \mid \forall \mathcal{B}' \subseteq \mathcal{A} : \underline{IE}_{C_a}(\mathcal{B}) \geq \underline{IE}_{C_a}(\mathcal{B}')\}$$

rules · Γ-maximin

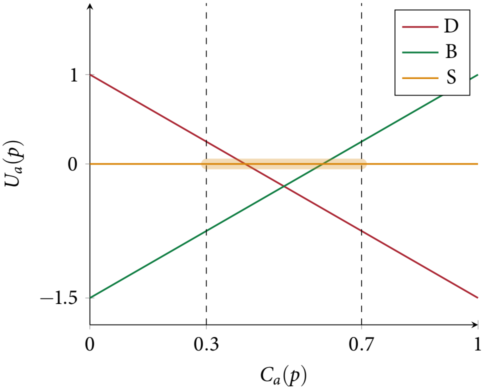

Illustration. A CONSERVATIVE risk profile with $[.3, .7]$ credence in p. The thicker, shaded lines mark the permitted solutions according to Γ-*maximin*, which only includes suspension.

rules · Γ-maximin

Good news: Γ-*maximin* is **unanimous**: If the members of your credal set agree on which belief sets maximize expected accuracy, these sets will be permitted by MAXIMIN.[11]

Bad news: Γ-*maximin* is **not strongly unanimous**: Unless the members of your credal set agree *exactly* on which options maximize expected accuracy, options that are not unanimously preferred may be permitted.

Γ-*maximin* is **not safe**: If there is disagreement among the members of your credal set, dominated choices may be permitted.

[11] Proof in Appendix.

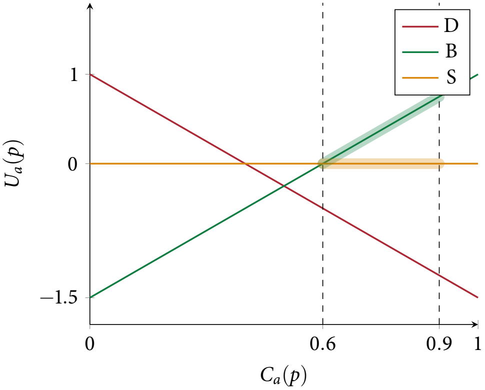

Counterexample. (Same as for *E-admissibility*.) A CONSERVATIVE risk profile with [.6, .9] credence in p. The thicker, shaded lines marks the permitted solutions according to Γ-*maximin*, including **belief** but also **suspension**, which is dominated and not unanimously preferred.

More bad news: Γ-*maximin* is **not persistent**.

If MAXIMIN permits the $\mathcal{B}$-solutions to all subproblems of an e.d.p. $\mathcal{A}_a$, it also permits $\mathcal{B}$ in $\mathcal{A}_a$:[12]

$$B \in \textsc{maximin}(\mathcal{A}_a) \Leftarrow B \cap A \in \textsc{maximin}(A_a) \text{ for all agendas } A \subseteq \mathcal{A}$$

But **the converse does not hold**: MAXIMIN may permit a belief set $\mathcal{B}$ in an e.d.p. $\mathcal{A}_a$, without permitting the $\mathcal{B}$-solutions to the subproblems of $\mathcal{A}_a$.

$$B \in \textsc{maximin}(\mathcal{A}_a) \not\Rightarrow B \cap A \in \textsc{maximin}(A_a) \text{ for all agendas } A \subseteq \mathcal{A}$$

[12] Proof in Appendix.

Counterexample. Consider an e.d.p. $\mathcal{A}_a = \{p, \neg p, q, \neg q\}$, where

- C_a contains two (probabilistic) credence functions $c_1, c_2 : c_1(p) = c_1(\neg q) = c_2(q) = c_2(\neg p) > .5$, their mixtures, and nothing more;
- U_a is NEUTRAL, with $V^S = 0$.

For the full e.d.p., we get that $\textsc{maximin}(\mathcal{A}_a) = \{\{p, q\}, \{\neg p, \neg q\}, \{\}\}$: these are the only options with a non-negative lowest *IE* (namely 0).

But $\textsc{maximin}(\{p, \neg p\}_a) = \{\{\}\}$ (and analogously for $\{q, \neg q\}_a$): $\underline{IE}_{C_a}(\{p\})$ must be the negative value given by $E_{c_2}(\{p\})$, and $\underline{IE}_{C_a}(\{\neg p\})$ the negative value given by $E_{c_1}(\{\neg p\})$. But $\underline{IE}_{C_a}(\{\}) = 0$, since $V^S = 0$.

Thus e.g., $\underline{\mathcal{B} = \{p, q\} \in \textsc{maximin}(\mathcal{A}_a)}$, yet $\underline{\mathcal{B} \text{ solves no subproblems of } \mathcal{A}_a}$.

In other words, Γ-*maximin* is not **c-consistent**: by contracting the agenda, rational agents can be forced to change their attitude towards propositions remaining from the original agenda.

In sum:

	e-admissibility	Γ-maximin
unanimity	✓	✓
strong unanimity	✗	✗
safety	✗	✗
e-consistency	✗	✓
c-consistency	✓	✗

If we agree that the rational agent avoids dominated choices, we can do one better by explicitly constraining rules to reject dominated choices:

***E-admissibility**.** Permissible choices maximize expected utility according to some member of your credal set, *and* are safe.

The characterizing choice function is $\textsc{e-max}^* := \textsc{e-max} \circ \textsc{non-dom}$, where $\textsc{non-dom}$ is the choice function that permits all and only safe options:

$$\begin{aligned}\textsc{non-dom}(\mathcal{A}_a) = \{\mathcal{B} \mid \neg\exists\mathcal{B}' \subseteq \mathcal{A} : &[\forall c \in C_a : E_c(\mathcal{B}') \geq E_c(\mathcal{B})]\ \& \\ &[\exists c \in C_a : E_c(\mathcal{B}') > E_c(\mathcal{B})]\}\end{aligned}$$

*E-admissibility** is also **strongly unanimous**: If the members of your credal set agree that some option maximizes expected accuracy, this option is permitted, and only such options are permitted: the others are collectively dominated.

For the same reason, the choice function MAXIMIN* := MAXIMIN ∘ NON-DOM is not only **safe** but also **strongly unanimous**.

Still, neither composite rule is **persistent**: E*-MAX fails *e-consistency* (consider the counterexample used for E-MAX), and MAXIMIN* fails *c-consistency* (consider the counterexample used for MAXIMIN).

One might hope that composing *E-admissibility* (which is only c-consistent) with Γ-*maximin* (only e-consistent) would yield a separable decision rule.

E-maximin. Permitted options maximize expected utility according to some member of your credal set, *and* maximize the lowest expected utility assigned by any member of your credal set.

The corresponding choice function is $\textsc{e-maximin} := \textsc{e-maximin} \circ \textsc{maximin}$.

Unfortunately, *E-maximin* is **not c-consistent**.

Counterexample. Consider an e.d.p. $\mathcal{A}_a = \{p, \neg p, q, \neg q\}$, where

- C_a contains two (probabilistic) credence functions $c_1, c_2 : c_1(p) = c_1(\neg q) = c_2(q) = c_2(\neg p) > .5$, their mixtures, and nothing more;
- U_a is NEUTRAL, with $V^S = 0$.

For the full e.d.p., we get that E-MAXIMIN$(\mathcal{A}_a) = \{\{p, q\}, \{\neg p, \neg q\}, \{\}\}$: these are the only options with a non-negative lowest IE (namely 0), and they each maximize for the linear average of c_1, c_2, c_x, which assigns both p and q a .5 credence.

But E-MAXIMIN$(\{p, \neg p\}_a) = \{\{\}\}$ (and analogously for $\{q, \neg q\}_a$): $\underline{IE}_{C_a}(\{\})$ must be the negative value given by $E_{c_2}(\{p\})$, and $\underline{IE}_{C_a}(\{\neg p\})$ the negative value given by $E_{c_1}(\{\neg p\})$. But $\underline{IE}_{C_a}(\{\}) = 0$, since $V^S = 0$. Again, $\{\}$ maximizes relative to c_x.

Thus $\underline{\mathcal{B} = \{p, q\} \in \text{E-MAXIMIN}(\mathcal{A}_a)}$, yet $\underline{\mathcal{B} \text{ solves no subproblems of } \mathcal{A}_a}$.

rules

	e-admissibility	Γ-maximin	e-admissbility* / Γ-maximin*	e-maximin
unanimity	✓	✓	✓	✓
strong unanimity	✗	✗	✓	✗
safety	✗	✗	✓	✗
e-consistency	✗	✓	✗/✓	✓
c-consistency	✓	✗	✓/✗	✗

In sum: Neither of these rules satisfy all of the suggested desiderata.

While each can be strengthened to guarantee **strong unanimity** and **safety**, **persistence**—specifically, *c-consistency*—is not as easily achieved.

Conclusion

conclusion

This talk introduced a basic formal framework for approaching the problem of establishing a Lockean thesis for imprecise credences: **imprecise EUT**.

With this in place, we considered a small selection of **decision rules** that might be used for determining the rationality of coarse-grained doxastic attitudes, without requiring that fine-grained doxastic attitudes be precise.

We suggested some basic **desiderata** for such rules, and found that neither rule considered had all the properties we were looking for.

Beside the rules discussed today, I have looked (in varying level of detail) on

- additional rules from practical imprecise decision theory, s.a. Walley 1991's *Maximality*,
- social choice rules on aggregating the preferences of the members of the credal set, s.a. classic *Majority voting* (ordinal) and *Utilitarianism* (cardinal).

Some issues: *Maximality* fails **persistence** (it is not e-consistent!)

Aggregation is not useful for imprecise choice generally, but hold more promise for the special case of credal sets in e.d.p.s, which yield *Arrow consistent* preference domains (given our assumptions).

Still, the social choice rules (i) require aggregation over **infinite domains**, and (ii) struggle with intuitive desiderata pertaining to **credal update**.

Persistence constrains how the coarse-grained attitudes of a rational agent change along with changes of her **agenda**, all else remaining fixed.

But there are also plausible constraints on how these attitudes change along with changes of an agent's **fine-grained attitudes**: i.e., when the agent *learns* things, updating her credal set.

(Strong) unanimity provides a weak constraint on coarse-grained attitude change in response to **precisification** of your credence in some proposition (weak because conditioned on agreement).

Our imprecise credences can also

- **shift**: be updated to include some new values *and exclude* some of the old,
- **dilate**: be updated to include new values *and preserve* the old.

conclusion

Spelling out additional desiderata + testing additional rules against them are part of the work in progress. **Input on any level is highly appreciated!**

Regarding both philosophical/conceptual and formal parts of the project:

- Is there anything particular about the work presented so far that **does not make sense** to you?
- Is there anything particular that you think **would make sense to develop in more detail**, or generally put more emphasis on?
- Is there anything additional that you think **would make sense to include discussion of**, make use of, or at least look up?

Thank you!

References

references

- Bradley, S. (2019). Imprecise Probabilities. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Spring 2019 edn., Metaphysics Research Lab, Stanford University.
- Dorst, K. (2019). Lockeans maximize expected accuracy, *Mind* 128(509), 175–211.
- Easwaran, K. (2016). Dr. Truthlove, or: How I learned to stop worrying and love Bayesian probabilities, *Nous* 50(4), 816–853.
- Ellsberg, D. (1962). 'Risk, Ambiguity, and the Savage Axioms, *Quarterly Journal of Economics*, 75(4): 643–669.
- Foley, R. (1992). The epistemology of belief and the epistemology of degrees of belief, *American Philosophical Quarterly* 29(2), 111–124.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of mathematical economics*.

references

- Greaves, H., & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607–632.
- Gärdenfors, P. and Sahlin, N.-E. (1982). Unreliable probabilities, risk taking, and decision making, *Synthese*, pp. 433–438.
- Hempel, C. (1962). Deductive-Nomological vs. Statistical Explanation. In H. Feigl, & G. Maxwell (Eds.) *Minnesota Studies in the Philosophy of Science*, vol. III, (pp. 98–169). Minneapolis: University of Minnesota Press.
- James, W. (1897). The Will to Believe. In *The Will to Believe, and other essays in popular philosophy*. New York: Longmans Green.
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making, *Philosophical perspectives* 24, 218–323.

references

- Keynes, J. M. (1921). *A treatise on probability*, Macmillan & Company.
- Levi, I. (1974). On Indeterminate probabilities, Journal of Philosophy, 71: 391–418.
- Leitgeb, H. and Pettigrew, R.: (2010a), An objective justification of Bayesianism I: Measuring inaccuracy, Philosophy of Science 77(2), 201–235.
- Leitgeb, H. and Pettigrew, R.: (2010b), An objective justification of Bayesianism II: The consequences of minimizing inaccuracy, Philosophy of Science 77(2), 236–272.
- Walley, P. (1991), *Statistical reasoning with imprecise probabilities*, Chapman & Hall.

Appendix: Proofs

proof · persistence of MAX

Let $\mathcal{A}_a$ be an arbitrary precise e.d.p.. We want to show that:
$\mathcal{B} \in \textsc{max}(\mathcal{A}_a) \Leftrightarrow [\text{for all } A^{\pm} \subset \mathcal{A}, \mathcal{B} \cap A^{\pm} \in \textsc{max}(A^{\pm}_a)]$.[13]

Assume that $\mathcal{B} \in \textsc{max}(\mathcal{A}_a)$. Where $c \in C_a$, this is equivalent to

$$\text{for all } \mathcal{B}' \subseteq \mathcal{A} : E_c(\mathcal{B}) \geq E_c(\mathcal{B}')$$

which the linearity of E lets us rewrite as

$$\text{for all } \mathcal{B}' \subseteq \mathcal{A} : \sum_{A^{\pm} \subset \mathcal{A}} E_c(\mathcal{B} \cap A^{\pm}) \geq \sum_{A^{\pm} \subset \mathcal{A}} E_c(\mathcal{B}' \cap A^{\pm}).$$

This can clearly hold if, and only if,

$$\text{for all } A^{\pm} \subset \mathcal{A} : \text{for all } \mathcal{B}' \subseteq \mathcal{A} : E_c(\mathcal{B} \cap A^{\pm}) \geq E_c(\mathcal{B}' \cap A^{\pm}),$$

which, by definition of MAX, says that $\mathcal{B} \cap A^{\pm} \in \textsc{max}(A^{\pm}_a)$ for all $A^{\pm} \subset \mathcal{A}$. $\square$

[13] I use the superscript $\pm$ to indicate that a set is *non-empty* and *closed under negation*.

proof · c-consistency of E-MAX

Let $\mathcal{A}_a$ be an arbitrary e.d.p.. We want to show that:
$\mathcal{B} \in \textsc{e-max}(\mathcal{A}_a) \Rightarrow [\text{for all } A^{\pm} \subset \mathcal{A}, \mathcal{B} \cap A^{\pm} \in \textsc{e-max}(A^{\pm}_a)]$.

Assume that $\mathcal{B} \in \textsc{e-max}(\mathcal{A}_a)$, or equivalently, that

$$\text{for some } c \in C_a\text{: for all } \mathcal{B}' \subseteq \mathcal{A}\text{: } E_c(\mathcal{B}) \geq E_c(\mathcal{B}').$$

By the linearity of E, we can rewrite this as

$$\text{for some } c \in C_a\text{: for all } \mathcal{B}' \subseteq \mathcal{A}\text{: } \sum_{A^{\pm} \subseteq \mathcal{A}} E_c(\mathcal{B} \cap A^{\pm}) \geq \sum_{A^{\pm} \subseteq \mathcal{A}} E_c(\mathcal{B}' \cap A^{\pm}).$$

which can hold if, and only if,

$$\text{for some } c \in C_a\text{: for all } \mathcal{B}' \subseteq \mathcal{A}\text{: for all } A^{\pm} \subseteq \mathcal{A}\text{: } E_c(\mathcal{B} \cap A^{\pm}) \geq E_c(\mathcal{B}' \cap A^{\pm}).$$

Rearranging the quantifiers yields the implied claim

$$\text{for all } A^{\pm} \subseteq \mathcal{A}\text{: for some } c \in C_a\text{: for all } \mathcal{B}' \subseteq \mathcal{A} : E_c(\mathcal{B} \cap A^{\pm}) \geq E_c(\mathcal{B}' \cap A^{\pm}).$$

This says that $A^{\pm} \subseteq \mathcal{A}$, there is some $c \in C_a$ such that for all $\mathcal{B}' \subseteq A^{\pm}$,
$E_c(\mathcal{B}) \geq E_c(\mathcal{B}')$—by definition, that $\mathcal{B} \in \textsc{e-max}(A^{\pm})$ for all $A^{\pm} \subset \mathcal{A}$. $\square$

proof · unanimity of MAXIMIN

Let $\mathcal{A}_a$ be an arbitrary molecular e.d.p. such that some $\mathcal{B} \subseteq \mathcal{A}$ maximizes expected utility given any $c \in C_a$. We want to show that $\mathcal{B} \in \textsc{maximin}(\mathcal{A}_a)$, or equivalently, that $\underline{IE}_{C_a}(\mathcal{B}) \geq \underline{IE}_{C_a}(\mathcal{B}')$ for all $\mathcal{B}' \subseteq \mathcal{A}$.

Consider $c \in C_a : E_c(\mathcal{B}) = \underline{IE}_{C_a}(\mathcal{B})$. By assumption, $\mathcal{B}$ is E-maximal for any member of C_a. In particular, $E_c(\mathcal{B}) \geq E_c(\mathcal{B}')$, for any $\mathcal{B}' \subseteq \mathcal{A}$. By definition of imprecise expected utility, this means that $E_c(\mathcal{B}) \geq \underline{IE}_{C_a}(\mathcal{B}')$, or equivalently, that $\underline{IE}_c(\mathcal{B}) \geq \underline{IE}_{C_a}(\mathcal{B}')$ for any $\mathcal{B}' \subseteq \mathcal{A}$. $\square$

proof · e-consistency of MAXIMIN

Let $\mathcal{A}_a$ be an arbitrary e.d.p.. We want to show that:
[For all $A^{\pm} \subset \mathcal{A}, \mathcal{B} \cap A^{\pm} \in \textsc{maximin}(A_a)] \Rightarrow \mathcal{B} \in \textsc{maximin}(\mathcal{A}_a)$.

Assume that, for all $A^{\pm} \subset \mathcal{A}$, $\mathcal{B} \cap A^{\pm} \in \textsc{maximin}(A_a^{\pm})$. That is:

$$\text{for all } A^{\pm} \subset \mathcal{A} \text{ and all } \mathcal{B}' \subseteq \mathcal{A} : \underline{IE}_{C_a}(\mathcal{B} \cap A^{\pm}) \geq \underline{IE}_{C_a}(\mathcal{B}' \cap A^{\pm}).$$

Consider $A^{\pm} \subset \mathcal{A} : \mathcal{A} - A^{\pm} = \{\varphi, \neg\varphi\}$ for some proposition φ. By assumption, $\underline{IE}_{C_a}(\mathcal{B} \cap A^{\pm}) \geq \underline{IE}_{C_a}(\mathcal{B}' \cap A^{\pm})$ and $\underline{IE}_{C_a}(\mathcal{B} \cap \{\varphi, \neg\varphi\}) \geq \underline{IE}_{C_a}(\mathcal{B}' \cap \{\varphi, \neg\varphi\})$. We must thus have that

$$\underline{IE}_{C_a}(\mathcal{B} \cap A^{\pm}) + \underline{IE}_{C_a}(\mathcal{B} \cap \{\varphi, \neg\varphi\}) \geq \underline{IE}_{C_a}(\mathcal{B}' \cap A^{\pm}) + \underline{IE}_{C_a}(\mathcal{B}' \cap \{\varphi, \neg\varphi\}),$$

which by definition of *IE* can be rewritten as

$$\underline{IE}_{C_a}((\mathcal{B} \cap (A^{\pm} \cup \{\varphi, \neg\varphi\})) \geq \underline{IE}_{C_a}(\mathcal{B}' \cap (A^{\pm} \cup \{\varphi, \neg\varphi\})).$$

By definition of $A^{\pm}$ and $\mathcal{B}, \mathcal{B}'$, we may rewrite this as $\underline{IE}_{C_a}(\mathcal{B}) \geq \underline{IE}_{C_a}(\mathcal{B}')$, which is equivalent to $\mathcal{B} \in \textsc{maximin}(\mathcal{A}_a)$, since $\mathcal{B}'$ was arbitrary. $\square$