Meaning and Artefact in the Measurement of Belief

Edward Elliott
CUNY, 2nd February 2024

Slides: tinyurl.com/EECUNY

Introduction I

- Beliefs come in degrees, which we can represent numerically
  - I'm **90% confident** that... (absolute values)
  - I'm **twice as confident** that... (ratio relations)

- Today's topic: the *measurement* of degrees of belief
  - **Roughly:** What do the numbers *mean*? What constitutes a *meaningful* difference between numerically distinct representations?
  - **Not really:** What *are* degrees of belief? (e.g., behaviourism or instrumentalism or representationalism or ...)
  - **Definitely not:** How do we *empirically determine* strengths of belief? (e.g., observation of betting behaviour)

Introduction II

The Measurement of Subjective Probability

Epistemic Accounts

Focus on the *internal* relational structure of the doxastic system — how doxastic states relate to other doxastic states

Decision-Theoretic Accounts

Focus on the *external* relational structure — how doxastic states relate to other psychological states (e.g., desires, intentions)

Notes on terminology

1. I might use the following interchangeably…
    - beliefs
    - graded beliefs
    - degrees of belief
    - subjective probabilities
    - credences
2. ‘Doxastic state’ = mental state with a belief-ish flavour
    - Not only degrees of belief
    - Comparative confidence
    - Judgements of evidential (in)dependence

Signposting

- Argument for an epistemic view (comparativism)
- The first problem — the shallow one
- The second problem — the deeper one
- Some stuff on measurement theory
- Lessons re: meaning, invariance, intrinsic structure

The Argument

Setup: decision-theoretic representations

A **decision-theoretic representation** of preferences consists in …

1. a representation of beliefs
2. a representation of desires
3. a decision rule to combine them

If preferences have **expected utility representation** $<pr, u>$ …

⇒ act α preferred to act β **iff** α has higher expected utility than β, relative to the probabilities *pr* and utilities *u*

$$\sum pr(S_i)u\big(\alpha[S_i]\big) > \sum pr(S_i)u\big(\beta[S_i]\big)$$

Setup: utility and interval-preserving transformations

Preferences have **expected utility representation** $<pr, u>$

$\Rightarrow$ α preferred to β iff $\sum pr(S_i)u\big(\alpha[S_i]\big) > \sum pr(S_i)u\big(\beta[S_i]\big)$

$\Updownarrow$ where $u^\star(x) = 9u(x) + 1$

Preferences have **expected utility representation** $< pr, u^\star >$

$\Rightarrow$ α preferred to β iff $\sum pr(S_i)u^\star\big(\alpha[S_i]\big) > \sum pr(S_i)u^\star\big(\beta[S_i]\big)$

Setup: temperature and interval-preserving transformations

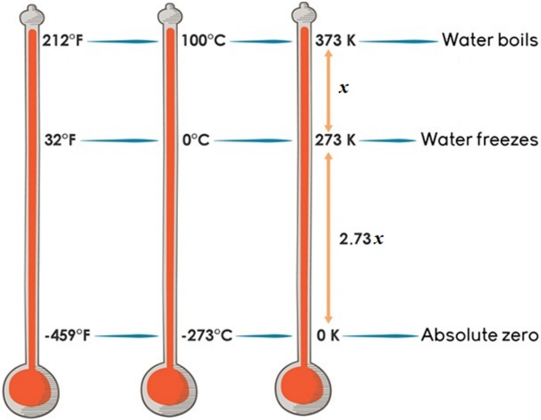

Setup: meaningfulness in utility

Step 1.

- An expected utility representation exists iff some other expected utility representation exists, with a different function for desires
- The alternative decision-theoretic representations of preference are all equally legitimate

∴ The alternative representations of desire they employ are also equally legitimate

Step 2.

- What's meaningful in a representation of desire is what it has in common with other equally legitimate representations

∴ There's no meaningful difference between u and $u^\star$

Zynda's argument: believabilities

Example (Zynda's original)

Take probability function *pr*, transform into **believability function**:

$$bel(S) \ = \ 9pr(S) + 1$$

- *bel* and *pr* have the same ordering
- *bel* takes values between 1 and 10
- $pr \to bel$ transformation does not preserve ratios
  i.e., if $pr(P) = r \cdot pr(Q)$, then $bel(P) \neq r \cdot bel(Q)$

Zynda's argument: equivalence of valuation maximising

Preferences have **expected utility representation** $<pr, u>$
$\Rightarrow$ α preferred to β iff $\sum pr(S_i)u(\alpha[S_i]) > \sum pr(S_i)u(\beta[S_i])$

$\Updownarrow$ where $bel(S) = 9pr(S) + 1$

Preferences have **valuation maximising representation** $<bel, u>$
$\Rightarrow$ α preferred to β iff

$$\sum bel(S_i)u(\alpha[S_i]) - u(\alpha[S_i]) > \sum bel(S_i)u(\beta[S_i]) - u(\beta[S_i])$$

Zynda's argument: proof of equivalence

$$pr(S) = \frac{bel(S) - 1}{9}$$

$$\sum pr(S_i)u\big(\alpha[S_i]\big) \;>\; \sum pr(S_i)u\big(\beta[S_i]\big)$$

substitute…

$$\sum \left[\frac{bel(S_i) - 1}{9}\right] u\big(\alpha[S_i]\big) \;>\; \sum \left[\frac{bel(S_i) - 1}{9}\right] u\big(\beta[S_i]\big)$$

drop constant…

$$\sum \big[bel(S_i) - 1\big] u\big(\alpha[S_i]\big) \;>\; \sum \big[bel(S_i) - 1\big] u\big(\beta[S_i]\big)$$

multiply out…

$$\sum bel(S_i)u\big(\alpha[S_i]\big) - u\big(\alpha[S_i]\big) \;>\; \sum bel(S_i)u\big(\beta[S_i]\big) - u\big(\beta[S_i]\big)$$

Zynda's argument: meaningfulness in belief

STEP 1.

- An expected utility representation exists iff some **valuation maxing** representation exists, with a different function for **belief**
- The alternative decision-theoretic representations of preference are all equally legitimate

∴ The alternative representations of belief they employ are also equally legitimate

STEP 2.

- What's meaningful in a representation of belief is what it has in common with other equally legitimate representations

∴ There's no meaningful difference between *pr* and *bel* (consequence: ratios in *pr* are **not** meaningful)

Zynda's argument: quotes

STEP 3. The **numerical ordering** is common to *pr* and *bel*...

> "*Now there are qualitative properties that subjective probabilities and any such "believability functions" would have to have in common [viz, the ordering] ... The concept of degree of belief on this strategy becomes a* **purely ordinal notion**"
>
> Zynda 2000, 64–5

> "*What* bel *and* pr *have in common is a comparative belief relation that ranks propositions according to an agent's confidence in their truth. In fact,* **all numerical degree of belief functions** *that figure in a [decision-theoretic representation] of the agent's preferences will represent the relation in question...*"
>
> Stefannson 2018, 382–3
> see also Frankish 2007, Stefansson 2016, Leitgeb 2021

Zynda's argument: comparativism

Comparativism: numerical probabilities are just a way of representing comparative confidence relations

$$P \succ Q \;\Leftrightarrow\; \text{more confidence in } P \text{ than } Q$$
$$P \sim Q \;\Leftrightarrow\; \text{same confidence in } P \text{ and } Q$$

pr and *bel* represent exactly the same comparative confidences, so there's no meaningful difference between them…

$$pr(P) > pr(Q) \quad \Leftrightarrow \quad P \succ Q \quad \Leftrightarrow \quad bel(P) > bel(Q)$$
$$pr(P) = pr(Q) \quad \Leftrightarrow \quad P \sim Q \quad \Leftrightarrow \quad bel(P) = bel(Q)$$

Summary

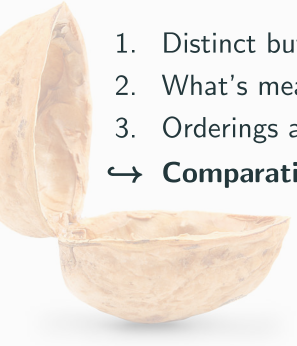

1. Distinct but equally legitimate representations of belief exist
2. What's meaningful is what's invariant between them
3. Orderings are invariant between them

↪ **Comparativism!**

Summary

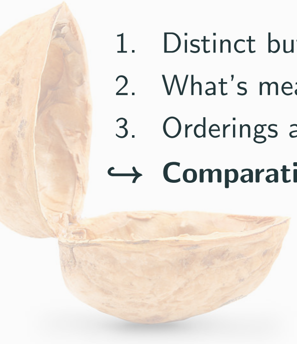

1. Distinct but equally legitimate representations of belief exist
2. What's meaningful is what's invariant between them
3. Orderings are invariant between them

$\hookrightarrow$ **Comparativism!**

Shallow problem: invalid inference
Deeper problem: premise 2 is false

The Shallow Problem

Linear transformations and difference ratios

$pr \rightarrow bel$ transformation preserves more than just the orderings — it also preserves **difference ratios**…

$$\frac{pr(P) - pr(Q)}{pr(R) - pr(S)} = \frac{bel(P) - bel(Q)}{bel(R) - bel(S)}$$

This is relevant because difference ratios

a) … matter for expected utility theory
b) … are not (generally) determined by comparative confidence

Example: difference ratios matter

Example (naïve comparativism & decision theory)

- $\alpha =$ bet for \$1 if P true, nothing otherwise
  $\beta =$ bet for \$2 if P is false, nothing otherwise
- Four propositions: $\Omega, P, \neg P, \varnothing$
- Confidence relation $\succsim$, with $\Omega \succ P \succ \neg P \succ \varnothing$
- pr represents $\succsim$ iff $1 > pr(P) > 0.5$
- α preferred to β iff $pr(\Omega) - pr(P) < pr(P) - pr(\neg P)$

Upshot: the role of numerical probability in expected utility theory requires more than just ordering information

Example: difference ratios matter

Example (naïve comparativism & decision theory)

- α = bet for \$1 if P true, nothing otherwise
  β = bet for \$2 if P is false, nothing otherwise
- Four propositions: $\Omega, P, \neg P, \varnothing$
- Confidence relation $\succsim$, with $\Omega \succ P \succ \neg P \succ \varnothing$
- *bel* represents $\succsim$ iff $10 > \mathit{bel}(P) > 5.5$
- α preferred to β iff $\mathit{bel}(\Omega) - \mathit{bel}(P) < \mathit{bel}(P) - \mathit{bel}(\neg P)$

Upshot: the role of numerical believability in val-maxing theory requires more than just ordering information

It’s the same everywhere else too…

Expected utility theory differentiates between ordinally-equivalent representations — presupposes meaningful extra-ordinal information

See also…

- other normative decision theories
- all major descriptive decision theories
- confirmation theory
- epistemic utility theory
- accuracy-first epistemology
- chance-credence principles
- information theory

Typical ‘epistemic’ responses

numerical probabilities represent something richer than a simple confidence ordering

impose more structure on comparative confidence

problems arise only in cases where there are many ordinally-equivalent representations of the confidence relation, but with further structural assumptions $\succsim$ will have a *unique* probabilistic representation

pluralism: posit additional doxastic relations to represent

if more than just ordinal information is meaningful, then perhaps the extra information corresponds to distinctive doxastic relations separate from comparative confidence (e.g., qualitative independence)

The Deeper Problem

The story so far

1. Alternative decision-theoretic representations generate distinct but equally legitimate numerical representations of belief
2. What's meaningful is only what's invariant across all the legitimate representations of belief
3. Difference ratios in probabilistic (and believabilistic) representations of belief seem to be meaningful

∴ Difference ratios *should* be invariant across all the other legitimate representations of belief... right?

Counterexample: schmelievabilities

Example (Schmynda variation)

Take probability *pr*, transform into **schmelievability function**:

$$sch(P) \;=\; bel(P)^2$$

- *pr* and *sch* have the same ordering
- *sch* takes values between 1 and 100
- *pr* $\rightarrow$ *sch* transformation does **not** preserve difference ratios

Counterexample: equivalence of schmaluation maximising

Preferences have **expected utility representation** $<pr, u>$

$\Rightarrow$ α preferred to β iff $\sum pr(S_i)u(\alpha[S_i]) > \sum pr(S_i)u(\beta[S_i])$

$\Updownarrow$ where $sch(S) = bel(S)^2$

Preferences have **schmaluation maximising representation** $<sch, u>$

$\Rightarrow$ α preferred to β iff

$$\sum \left[\sqrt{sch(S_i)} - 1\right] u(\alpha[S_i]) > \sum \left[\sqrt{sch(S_i)} - 1\right] u(\beta[S_i])$$

Counterexample: proof of equivalence

$$pr(S_i) = \frac{bel(S_i) - 1}{9}, \quad bel(S_i) = \sqrt{sch(S_i)}$$

$$\sum pr(S_i) u\big(\alpha[S_i]\big) \;\geq\; \sum pr(S_i) u\big(\beta[S_i]\big)$$

substitute…

$$\sum \left[\frac{bel(S_i) - 1}{9}\right] u\big(\alpha[S_i]\big) \;\geq\; \sum \left[\frac{bel(S_i) - 1}{9}\right] u\big(\beta[S_i]\big)$$

drop constant…

$$\sum \big[bel(S_i) - 1\big] u\big(\alpha[S_i]\big) \;\geq\; \sum \big[bel(S_i) - 1\big] u\big(\beta[S_i]\big)$$

substitute…

$$\sum \big[\sqrt{sch(S_i)} - 1\big] u\big(\alpha[S_i]\big) \;\geq\; \sum \big[\sqrt{sch(S_i)} - 1\big] u\big(\beta[S_i]\big)$$

It's (almost) all meaningless?

1. The equivalence proofs rely only on the invertibility of the transform — generalises easily to *any* bijective automorphism
2. This includes transformations that do not preserve ratios, or ratios of differences, *or even orderings*…
3. (Almost) nothing is invariant across **all** the representations of belief for all possible decision-theoretic representations

What’s going on?

Difference ratios in probabilistic representations are meaningful

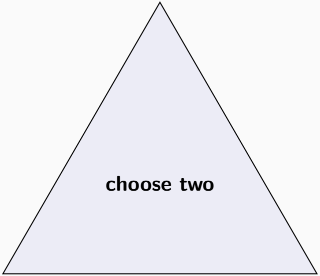

What’s meaningful is what’s invariant across all legitimate representations of belief

Alternative decision-theoretic representations provide legitimate representations of belief

Possible responses

1. **Bite the bullet:** what's meaningful is what's invariant across the alternative representations, which is (almost) nothing

2. **Deny legitimacy:** the alternative *non*-expected utility representations were not so legitimate after all

3. **Rethink invariance:** correct to see invariance as indicative of meaningfulness, but this needs to be understood correctly

What's going on?

Difference ratios in probabilistic representations are meaningful

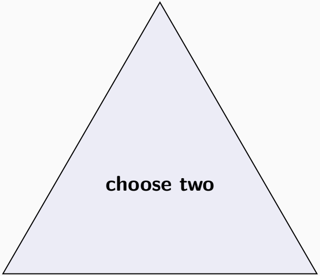

What's meaningful is what's invariant across all legitimate representations of belief

Alternative decision-theoretic representations provide legitimate representations of belief

Meaning and Invariance

Length ratios

Spot is twice as long as Fluffy

Harry is two-thirds as long as Fluffy

Additive measures of length

The system of lengths

$\mathbf{O} = \{o_1, o_2, \dots\}$ is a set of **objects**
$\succsim$ is the **at least as long as** relation
$\oplus$ is the **concatenation** operation

$$< \mathbf{O}, \succsim, \oplus >$$
$$\Downarrow \varphi$$
$$< \mathbb{R}^{>0}, \geq, + >$$

φ maps $\mathbf{O}$ into $\mathbb{R}^{>0}$ such that:
$o_1 \succsim o_2$ iff $\varphi(o_1) \geq \varphi(o_2)$
$o_1 \oplus o_2 = o_3$ iff $\varphi(o_1) + \varphi(o_2) = \varphi(o_3)$

The additive representation

$\mathbb{R}^{>0}$ is the set of **positive reals**
$\geq$ is the **at least as great as** relation
$+$ is the **addition** operation

Internal structure and length

In the case of **length**…

1. The numbers represent the internal relational structure of the "system of lengths" $< \mathbf{O}, \succsim, \oplus >$ — roughly: how lengths relate to other lengths
2. What the numbers *mean* can be fully explained without reference to how length relates to other quantities
3. Since ratios are *meaningful* (for additive measures), that meaning must be reflected in the internal structure of the length system

Additive measures and ratio-preserving transformations

Ratios are always invariant across additive measures of length

$$< \mathbf{O}, \succsim, \oplus > \quad \Leftrightarrow \quad < \mathbf{O}, \succsim, \oplus >$$

$$\Downarrow \varphi \qquad\qquad \Downarrow \psi$$

$$< \mathbb{R}^{>0}, \geq, + > \qquad < \mathbb{R}^{>0}, \geq, + >$$

if $\varphi \to \psi$ is ratio-preserving

e.g., *meters*, *inches*, *lightyears* — different units, same ratios

But…

ratio relations are meaningful (for additive measures)

$\neq$

ratio relations are invariant across *all* possible measures

Multiplicative measures of length

The system of lengths

$\mathbf{O} = \{o_1, o_2, \ldots\}$ is a set of **objects**
$\succsim$ is the **at least as long as** relation
$\oplus$ is the **concatenation** operation

$$< \mathbf{O}, \succsim, \oplus >$$
$$\Downarrow \varphi$$
$$< \mathbb{R}^{>1}, \geq, \times >$$

φ maps $\mathbf{O}$ into $\mathbb{R}^{>1}$ such that:
$o_1 \succsim o_2$ iff $\varphi(o_1) \geq \varphi(o_2)$
$o_1 \oplus o_2 = o_3$ iff $\varphi(o_1) \times \varphi(o_2) = \varphi(o_3)$

The multiplicative representation

$\mathbb{R}^{>1}$ is the set of **reals** $>$ **1**
$\geq$ is the **at least as great as** relation
$\times$ is the **multiplication** operation

Meaning is relative

- Ratios are not preserved between any multiplicative measures, nor between any additive and multiplicative measures
- There are infinite alternative numerical structures in which the length system $< \mathbf{O}, \succsim, \oplus >$ can be represented
- (Almost) nothing is invariant across **all** possible numerical representations of length

Meaning is relative

Lesson 1. meaningfulness only sensibly defined *given* a fixed choice of representational structure — what's meaningful is what's invariant relative to a 'representational format'[1]

[1] cf. Pfzangl 1968, Narens 1980, Roberts 1981, Luce et al. 1990...

Back to Zynda

Difference ratios in probabilistic representations are meaningful

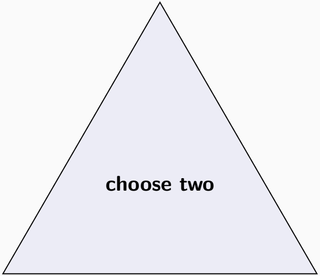

What's meaningful is what's invariant across all legitimate representations of belief

Alternative decision-theoretic representations provide legitimate representations of belief

Meaning and Intrinsic Structure

Epistemic theories of belief measurement

According to **epistemic theories**…

1. The numbers represent the intrinsic relational structure of the doxastic system — roughly: how doxastic states relate to other doxastic states
2. What the numbers *mean* can be fully explained without reference to how the doxastic system relates to other psychological phenomena
3. If any extra-ordinal information is *meaningful*, then that fact must be reflected in the internal structure of the doxastic system

Intrinsic and non-intrinsic structure

Option 1: a numerical model of belief represents (only) the internal qualitative structure of the doxastic system

Option 2: a numerical model of belief represents (also) something about how the doxastic system relates to other parts of our psychology

Conjoint measurement

Conjoint measurement involves the simultaneous representation of several related quantities

- Typical application: quantities **A** and **B** jointly determine quantity **C**
- Representation of **C** has three parts:
  - representation of **A**
  - representation of **B**
  - combination rule R

Example 1. discomfort determined by humidity and temperature, $dis = hum + temp$

Example 2. density determined by mass and volume, $den = \frac{mass}{vol}$

Additive conjoint measurement I

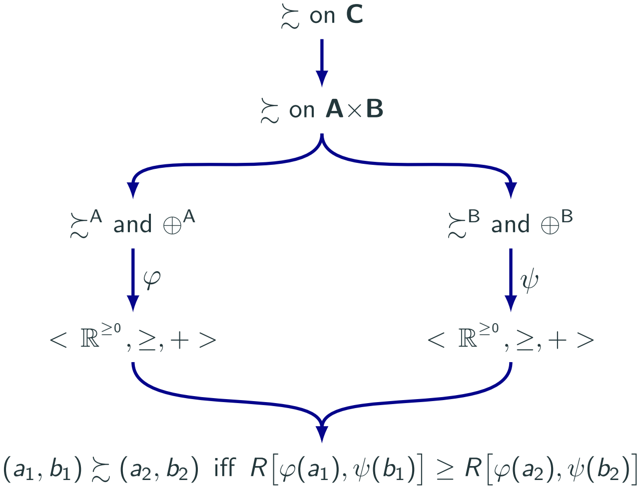

Additive conjoint measurement II

$\Rightarrow$ the representation of **A** depends on its relations to **B** and **C**

1. Constructing $\succsim^{A}$

   $\approx$ $a_2 \succ^{A} a_1$ iff changing from a_1 to a_2 (holding the level of **B** fixed) will increase the level of **C**

2. Constructing $\oplus^{A}$

   $\approx$ $a_1 \oplus a_2 = a_3$ iff changing from a_1 to a_3 (given a specified level of **B**) has the same effect on **C** as a_1 combined with an increase in **B** equal in effect on **C** to a_2

Upshot: in conjoint systems, we cannot fully explain what the numbers *mean* in measurement of **A** without referencing its relations to **B** and **C**

Lessons

Lesson 2. decision-theoretic representations are conjoint — the decision rule is part of the 'representational format'

- comparing **expected utility** and **val-maxing** representations is like comparing additive and multiplicative measures of length
- i.e., what's *invariant* between *pr* and *bel* (and *sch* and...) is *irrelevant* to what's meaningful in those functions

Lessons

Lesson 3. Zynda-style examples actually show that (difference) ratios in *pr* *are* meaningful for expected utility representations

- if the transformation of the belief function doesn't preserve ratios, then the decision rule must be altered
- so if an alternate representation uses the same decision rule, the belief function must have same ratios in common with *pr*
- compare desires: no meaningful difference between u and $u^\star$ *because* $u \to u^\star$ transformation requires no change in decision rule

Invariance in decision-theoretic representations

expected utility / valuation / schmaluation representations capture the same underlying system of relationships between…

1. beliefs about states
2. desires for outcomes
3. preferences over acts

$$\left.\begin{array}{l} \sum pr(S_i)\,u\big(\alpha[S_i]\big) \\[2em] \sum \left[\frac{bel(S_i)-1}{9}\right] u\big(\alpha[S_i]\big) \\[2em] \sum \left[\sqrt{sch(S_i)}-1\right] u\big(\alpha[S_i]\big) \end{array}\right\} \quad \sum f(S_i)\cdot g\big(\alpha[S_i]\big)$$

Invariance in conjoint systems

If $\boldsymbol{F}$ in Newtons, $\boldsymbol{m}$ in kilograms, $\boldsymbol{a}$ in meters/second^2

$$\boldsymbol{F} = \boldsymbol{m}\boldsymbol{a}$$

If $\boldsymbol{F}$ in Newtons, $\boldsymbol{m}$ in grams, $\boldsymbol{a}$ in meters/second^2

$$\boldsymbol{F} = \frac{\boldsymbol{m}\boldsymbol{a}}{1000}$$

If $\boldsymbol{F}$ in nega-Newtons, $\boldsymbol{m}$ in lbs, $\boldsymbol{a}$ in multiplicative-meters/second^2

$$\boldsymbol{F} = \frac{-5\boldsymbol{m}\log_2(\boldsymbol{a})}{11}$$

Conclusion

If alternate decision-theoretic representations provide us with legitimate alternate representations of belief…

1. meaningfulness is not found in what's invariant across the different representations of belief (considered in isolation)
2. understanding what's meaningful may require understanding of how the doxastic system relates to other psychological systems

Appendix

Appendix: utilities and external structure

extra-ordinal information for **desires** is not explained via *internal* structure of the desire system in isolation, but in terms of how desires relate to **beliefs**…

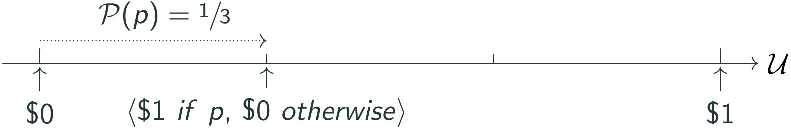

Appendix: utilities and external structure

extra-ordinal information for **desires** is not explained via *internal* structure of the desire system in isolation, but in terms of how desires relate to **beliefs**…

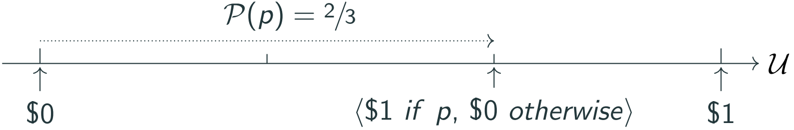

Varieties of epistemic theory

The system of beliefs

$\mathbf{P} = \{P, Q, \ldots\}$ is set of **propositions**

$\succsim$ is the **confidence** relation

$$< \mathbf{P}, \succsim >$$
$$\Downarrow \varphi$$
$$< \mathbb{R}^{[0,1]}, \geq >$$

φ maps $\mathbf{P}$ into $\mathbb{R}^{[0,1]}$ such that:
$P \succsim Q$ iff $\varphi(P) \geq \varphi(Q)$

The probabilistic representation

$\mathbb{R}^{[0,1]}$ is the **[0,1] interval of reals**

$\geq$ is the **at least as great as** relation

Varieties of epistemic theory

The system of beliefs

$\mathbf{P} = \{P, Q, \ldots\}$ is set of **propositions**
$\succsim$ is the **confidence** relation
$\sqcup$ is the **disjoint unions** operation

$$< \mathbf{P}, \succsim, \sqcup >$$
$$\Downarrow \varphi$$
$$< \mathbb{R}^{[0,1]}, \geq, + >$$

φ maps $\mathbf{P}$ into $\mathbb{R}^{[0,1]}$ such that:
$P \succsim Q$ iff $\varphi(P) \geq \varphi(Q)$
$P \sqcup Q = R$ iff $\varphi(P) + \varphi(Q) = \varphi(R)$

The probabilistic representation

$\mathbb{R}^{[0,1]}$ is the **[0,1] interval of reals**
$\geq$ is the **at least as great as** relation
$+$ is the **addition** operation

Varieties of epistemic theory

The system of beliefs

$\mathbf{P} = \{P, Q, \dots\}$ is set of **propositions**
$\succsim$ is the **confidence** relation
$\sqcup$ is the **disjoint unions** operation
$\perp$ is a **qual. independence** relation

$$< \mathbf{P}, \succsim, \sqcup, \perp >$$

$$\Downarrow \varphi$$

$$< \mathbb{R}^{[0,1]}, \geq, +, \times >$$

φ maps **P** into $\mathbb{R}^{[0,1]}$ such that:
$P \succsim Q$ iff $\varphi(P) \geq \varphi(Q)$
$P \sqcup Q = R$ iff $\varphi(P) + \varphi(Q) = \varphi(R)$
$P \perp Q$ iff $\varphi(P) \times \varphi(Q) = \varphi(P \cap Q)$

The probabilistic representation

$\mathbb{R}^{[0,1]}$ is the **[0,1] interval of reals**
$\geq$ is the **at least as great as** relation
$+$ is the **addition** operation
$\times$ is the **multiplication** operation