Decision Under Normative Uncertainty

Franz Dietrich
PSE & CES & CNRS

Brian Jabarian
U. Paris 1 & PSE

Second Workshop on Coping With Heterogeneous Opinions

Paris School of Economics

29 November 2018

Empirical vs. normative uncertainty

- Classical empirical uncertainty: uncertainty about empirical facts.
  - *Ex*: Does a medical treatment cure the patient? What are the side effects?
- Normative uncertainty: uncertainty about value facts.
  - *Ex*: Is curing the patient worth the side effects? How much does the patient's will count? What is the correct inequality aversion?
  - *More generally*: What is the correct normative theory? (Is it utilitarianism, some egalitarianism, some prioritarianism, some deontology, ...?)

Should we close down nuclear plants?

Two dimensions of this debate:

- *empirical uncertainty*: Will there be earth quakes? human errors? technological progress? etc.
- *normative uncertainty*: How evaluate burdens for future generations? What is the correct intergenerational discounting factor? How trade off between quality of life and probability of death in accidents? etc.

Goal: incorporate normative uncertainty into decision models

Why important?

Understanding both sides of (social and internal) deliberation

‘Value’ could stand for...

- individual well-being,
- social welfare,
- moral value,
- legal value,
- artistic value,
- ...

Conceptualizing normative uncertainty within Savage's framework

Coming from Savage's decision theory, one might think of

- empirical uncertainty as uncertainty about the nature state (interpreted as the empirical state of the world)
- normative uncertainty as uncertainty about the value/utility of consequences.

Classical EU-agents have only empirical uncertainty: they do not know the state, but know ('have') exact utilities of consequences.

Note our cognitive re-interpretation of ‘utility’

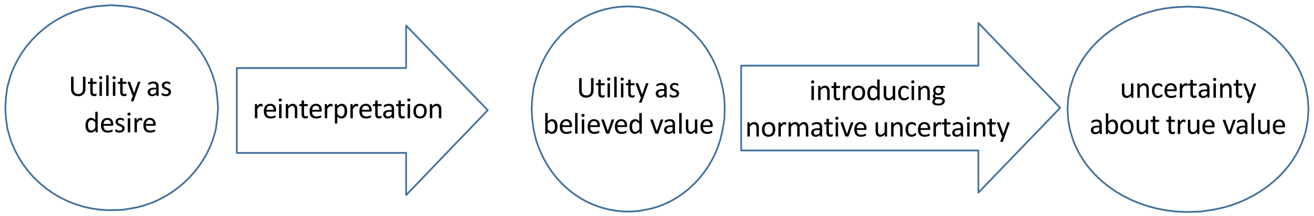

Figure 1: In 2 steps in normative uncertainty

From a Humean belief/desire model to a cognitivist model

Normative uncertainty: philosophically meaningful?

- Normative uncertainty presupposes (beliefs about) normative facts.
- ‘Normative facts’?? Don’t worry: these facts can be objective or subjective, universal or relative, ...
  I’ll spare you with philosophical debates around ‘facts’.

Normative uncertainty:

formally different?

- A legitimate question! (Which I had too, 1 year ago.)
- Modelling normative uncertainty as ordinary choice-theoretic uncertainty fails.
- So: normative uncertainty differs not just interpretively, but also formally.

Philosophers have started formal work on normative uncertainty

- MacAskill (2014, 2016), Greaves & Ord (2018), Lockhart (2000), Ross (2006), Sepielli (2009), Barry & Tomlin (2016)
- Some points of focus:
  - cardinal vs ordinal value
  - comparable vs non-comparable value
  - individual vs collective choice
  - consequentialist vs non-consequentialist evaluations

The Question

- How evaluate options under normative uncertainty?
  –> What's the 'meta-value' under uncertainty about '1^{st}-order value'?

Plan

1. The classical ‘expected-value theory’
2. An alternative ‘impartial value theory’

Options and Valuations

Consider:

- a set A of '*options*', the objects of evaluation
  - choice options, policy measures, social arrangements, income distributions, ...
  - (For now we set aside empirical uncertainty. But in principle options could contain empirical uncertainty.)

Valuations

- a finite set $\mathcal{V}$ of '*valuations*' v, assigning to each option $a \in A$ its *value* $v(a)$ in $\mathbb{R}$.
  - They might represent rival normative theories, normative intuitions, value judgments, 'social preferences', ...
  - $\mathcal{V}$ might consist of:
    - a utilitarian and a Rawlsian valuation, or
    - 'similar' valuations differing in a parameter, e.g., in a discounting factor, or inequality-aversion degree, or prioritarian degree, ...

Value versus vNM utility

Beliefs about value

Consider further:

- a probability function Pr assigning to each valuation v in $\mathcal{V}$ its subjective correctness probability $Pr(v) \geq 0$, where $\sum_{v \in \mathcal{V}} Pr(v) = 1$.

Meta-theories

- What is the *overall* value of each option, given one's normative uncertainty?
- An answer is a *'meta-'valuation*, assigning to each option in A its 'overall' value.
- Prominent proposal: the **expected-value theory** 'EV' which valuates each option $a \in A$ by its expected value:

$$EV(a) = \sum_{v \in \mathcal{V}} Pr(v)v(a).$$

EV is neutral to normative risk

Neutrality to normative risk is implausible if aversion to empirical risk is certainly correct

What does it mean that aversion to empirical risk is certainly correct?

- Assume options in A contain empirical uncertainty. say they are vNM lotteries on a set X of 'outcomes'.
- The value of an outcome x in X is the value of the sure lottery in A which yields x.
- The risk attititude of a valuation $v \in \mathcal{V}$ is given by how $v(a)$ compares to the expected outcome-value $\sum_{x \in X} a(x)v(x)$.
- Risk-aversion is certainty correct if $v(a) < \sum_{x \in X} a(x)v(x)$ for all non-sure lotteries a and all $v \in \mathcal{V}$ s.t. $Pr(v) \neq 0$.

The attitude of *EV* to empirical risk is *impartial*: it is guided by the risk-attitudinal beliefs

- EV is neutral (averse, prone) to *empirical* risk if all $v \in \mathcal{V}$ of non-zero correctness probability $Pr(v)$ are risk-neutral (-averse, -prone). Formally, EV evaluates options without normative risk at (below, above) the option's expected outcome value if each $v \in \mathcal{V}$ s.t. $Pr(v) \neq 0$ does so.
- 'Impartiality' of risk attotides can be defined precisely.

In the paper we define 3 alternatives to EV, with different risk attitudes

	neutral to nor. risk	impartial to nor. risk
neutral to emp. risk	*'fully expectational value'*	*'dual expected value'*
impartial to emp. risk	*'expected value'*	*'impartial value'*

Our favourite: the impartial value theory.

How is it defined?

Value prospects

- A **value prospect** is a lottery over value levels in $\mathbb{R}$.
- Each option $a \in A$ generates two types of value prospect, depending on whether we consider just empirical or also normative uncertainty:
  - a's **value prospect under** $v \in \mathcal{V}$ is denoted $p_{a,v}$ and given by:

$$\begin{aligned} p_{a,v}(k) &= \text{ prob. of an outcome of value } k \text{ under } v \\ &= \sum_{x \in X : v(x) = k} a(x). \end{aligned}$$

  - a's **value prospect** simpliciter is denoted p_a and given by:

$$\begin{aligned} p_a(k) &= \text{ prob. of an outcome of value } k \\ &= \sum_{(v,x) \in \mathcal{V} \times X : v(x) = k} \underbrace{Pr(v)a(x)}_{\text{prob. of } (v,x)}. \end{aligned}$$

Impartial Value defined

- Each valuation v in $\mathcal{V}$ can be taken to evaluate not just options a, but also value prospects p:[1]

  $v(p) =$ value $v(a)$ of options a with value prospect $p_{a,v} = p$.

- The **impartial theory** 'IV' evaluates each option $a \in A$ by the expected evaluation of its value prospect:

$$IV(a) = \sum_{v \in \mathcal{V}} Pr(v) v(p_a).$$

[1] This definition presupposes a technical assumption: for each valuation v in $\mathcal{V}$ and value prospect p, let there exist a corresponding option a in A whose value prospect $p_{a,v}$ is p, and moreover let any two such options a in A have same value $v(a)$.

IV versus *EV*

- Assume that being risk-averse is certainly correct, i.e., only risk-averse theories in $\mathcal{V}$ have positive probability.
- The expected value $EV(a) = \sum_{v \in \mathcal{V}} Pr(v)v(a)$ contains a risk premium for *empirical* risk, because each '$v(a)$' contains a premium for the (*empirical*) risk in a.
- The impartial value $IV(a) = \sum_{v \in \mathcal{V}} Pr(v)v(p_a)$ contains a risk premium for *empirical and normative* risk, because each '$v(p_a)$' contains a premium for the (*empirical and normative*) risk in p_a.

Ex-ante vs. ex-post approach

- Famous question in ethics and aggregation theory: should competing evaluations of uncertain prospects be aggregated before or after resolution of uncertainty? (See, e.g., Fleurbaey 2010, Fleurbaey and Zuber 2017.)
- We have two types of uncertainty, so four approaches:

	normatively ex-post	normatively ex-ante
empirically ex-post	fully expectational value	dual expected value
empirically ex-ante	expected value	impartial value

Why do we base *IV* on an expectation?

- Is IV not risk-neutral through the back door, through taking the *expectation* of the $v(p_a)$ ($v \in \mathcal{V}$)?
- No, because each $v(p_a)$ ($v \in \mathcal{V}$) already contains a premium for all the risk in the option a, empirical and normative. Defining $IV(a)$ as a value below that expectation would amount to a 'double risk premium'.