

Normative Uncertainty II

January 23, 2018

Maximize InterTheoretic Expected Value (MITE)

What should you do in the face of normative uncertainty?

Answer: "Do whatever maximizes expected intertheoretic moral value."

$$EIMV(\phi) = \sum_i Cr(T_i) \cdot V_{T_i}(\phi)$$

Against Moral Hedging

MITE is a form of moral hedging. Ittay Nissan-Rosen argues against moral hedging in general.

Moral Hedging Thesis: "[W]hen choosing under conditions of moral uncertainty, a morally motivated rational agent ought to take into account not only his degrees of belief in different moral theories, but also the degrees of moral value these theories assign to the acts he can choose from." [p. 350]

Claim 1: A morally motivated rational agent can be uncertain about the true moral theory's attitude toward risk. (This assumes that moral theories, in addition to saying what we *objectively* ought to do, tell us what we *subjectively* ought to do under conditions of non-normative uncertainty.)*

Claim 2: If you assign positive credence to two moral theories that differ in their attitudes toward risk, there is a conflict between the Moral Hedging Thesis and the following two dominance principles:

- o *The Moral Dominance Assumption:* If, according to all moral theories you assign positive credence, ϕ has more moral value than ψ , then ϕ is morally better than ψ .
- o *Standard Dominance:* If, for all ways the world might be S , $(\phi \wedge S)$ is morally better than $(\psi \wedge S)$, then ϕ is morally better than ψ .

Assume you are equally confident in two moral theories, T_1 and T_2 , that take different attitudes towards risk. You face the following moral decision problem (under normative & non-normative uncertainty):

Main Question: What ought you — a morally conscientious and minimally rational person — to do when you are uncertain about what you, morally, ought to do?

The Problem of Intertheoretic Value Comparisons.

1. *No Common Moral Currency.* Among moral theories that provide complete cardinal rankings of all options, there's no obvious way to *compare* these moral values *across* different moral theories. [Analogous to The Problem of Interpersonal Comparisons of Utility, but worse]
2. *Non-EVM-Representable Theories.* Not all moral theories provide complete cardinal rankings of your options. And it's not obvious that we can represent them in such a way as implicitly doing so.
 - (a) *Supererogatory Actions.* Moral theories that allow for supererogation cannot be represented with a value-function because they violate the Completeness Axiom.
 - (b) *Absolutist Moral Theories.* Moral theories that say, e.g., "Never lie!" or "Never murder!" either violate the Continuity Axiom or (if we give that up) will assign infinite values that threaten to swamp our intertheoretic expected moral value calculations.

*Could this be denied? "[I]f there is no 'right' attitude toward risk to adopt when making moral choices (if it is not 'immoral' to adopt any attitude toward risk with respect to moral value), nothing constrains the set of permissible choices given the agent's uncertainty. The moral hedging thesis turns out to be empty under this approach." [p. 355]

A Moral Decision Problem

	HEADS	TAILS		
<i>a</i>	A	A		
<i>b</i>	B	B		
<i>c</i>	C	C		
<i>ℓ</i>	A	B		
	A	C	B	
T_1 :	10	3	0	2.5 (<i>risk-averse</i>)
T_2 :	0	6	10	5 (<i>risk-neutral</i>)

	HEADS		TAILS		Expected Moral Values	
	T ₁	T ₂	T ₁	T ₂	EMV _{T₁}	EMV _{T₂}
<i>a</i>	10	0	10	0	10	0
<i>b</i>	0	10	0	10	0	10
<i>c</i>	3	6	3	6	3	6
<i>ℓ</i>	10	0	0	10	2.5	5

If the Moral Hedging Thesis is correct, then, according to The Moral Dominance Assumption, *c* is morally better than *ℓ*.

However, if you're risk-neutral with respect to normative uncertainty, then, according to Standard Dominance, *ℓ* is morally better than *ψ*.**

There appears to be a conflict here between moral hedging and rationality. Should we, then, reject the Moral Hedging Thesis? Are there other ways to respond to this problem?

Normative Externalism

As we saw last week, the subjective 'ought' is meant to play two theoretical roles: (1) it is *action-guiding* and (2) it tracks *praise-* and *blame-worthiness*.

Can it play these roles in the face of *normative* uncertainty?

- *Weatherson*: The only moral decision-rule is "Do what's morally right!"
 1. Hedging against moral risk cannot make an otherwise morally permissible action impermissible.
 2. Moral hedging reveals a character flaw: you are caring about what's morally right *de dicto*, not *de re*, which is fetishistic.
- *Harman*: Moral ignorance doesn't exculpate (in the way that non-normative ignorance sometimes does).

Regress Problem: As we've seen, it's not obvious what to do in the face of normative uncertainty; we don't know what the correct *metanormative* theory is.

So don't we also need to determine what to do when you don't know what you should because of normative uncertainty? That is, don't we need a *metametanormative* theory? ... And why think we can stop there?

We should stop before the regress gets going: *you should do what you morally ought to do*.

**Here's why. If HEADS, then

$$\begin{aligned}
 V(\ell) &= \frac{1}{2} \cdot V_{T_1}(A) + \frac{1}{2} \cdot V_{T_2}(A) \\
 &= \frac{1}{2} \cdot (10 + 0) = 5 \\
 V(c) &= \frac{1}{2} \cdot V_{T_1}(C) + \frac{1}{2} \cdot V_{T_2}(C) \\
 &= \frac{1}{2} \cdot (3 + 6) = 4.5
 \end{aligned}$$

Likewise, for TAILS. So, for all the ways the world might be, *ℓ* is morally better than *c*.

If moral ignorance doesn't exculpate, and there is a regress problem (and where are our credences in moral theories supposed to be coming from anyway?), it doesn't look like the subjective 'ought' is able to play the role it was meant to play.

If we accept Normative Externalism about normative uncertainty, then should we also accept it about non-normative uncertainty?