Lec. 1: A Puzzle about Rational Revisability

An incompatible foursome:

1. At any time, a person possesses a "highest level epistemic norm", that constitutes the person's standards of rational formation and retention of beliefs at that time.

2. (Assuming 1) It's not possible for the person to rationally revise that highest level epistemic norm under any conditions.

3. Any sufficiently high level rational norm must include a logic (powerful enough to be of use).

4. For any logic (that is powerful enough to be of any use), it would be possible for the person to rationally revise that logic under certain conditions.

Claim 1: At any time, a person possesses a "highest level epistemic norm", that constitutes the person's standards of rational formation and retention of beliefs at that time.

What does 'highest level epistemic norm' mean?

Indeed, what does 'epistemic norm' mean?

Take an epistemic norm to be a policy: a policy both

(i) for *believing* (or believing to a certain degree)

and (ii) for *acting so as to improve one's epistemic situation,* e.g. by

trying to gather more evidence,

or to think up more possible explanations,

or to determine whether an answer to a question of interest follows from things one already accepts.

Tempting to say that it's a policy *to which an agent has some kind of deep commitment*. But I want to be able to discuss the norm in abstraction from agents, so I won't build this in.

An alternative way of conceiving epistemic norms: as *normative claims*. (E.g.: "You *shouldn't* believe a conjunction without believing the conjuncts".)

But I want to think of the normative claims as generated by *policies to which the agent is committed*.

> We have a policy of not believing a conjunction without believing the conjuncts; this policy plays a very central role in our thinking.
>
> When we say that we *shouldn't* believe a conjunction without believing the conjuncts, we are expressing our commitment to this policy.

Focusing on policies rather than on normative claims is intended to leave open a kind of normative anti-realism, which I'll discuss in Lecture 5.

But my focus is *compatible* with normative realism.

The policy of not believing a conjunction without believing the conjuncts clearly has a status that contrary policies don't have. According to the normative realist, this special status is: being objectively correct. That is, following it involves *believing as one objectively should believe*.

I'm skeptical, but don't want to pre-judge the issue. (While it is *prima facie* plausible in the case of deductive policies like the above, it is much less so for, e.g., inductive policies. And I'll eventually argue it to be wrong even in the deductive case.)

Speaking of norms as policies doesn't pre-judge the issue. The issue of realism is the issue of whether it makes sense to talk of some norms (policies) as being "objectively correct", or "the objectively right ones to employ".

An anti-realist thinks that makes no sense, but can still talk about

- which norms a person does follow
- the advantages and disadvantages that certain norms have over others.

A “low level” epistemic norm:

> believe what one reads in the NY Times, unless it appears under the byline of Elizabeth Bumiller.

This is a policy that I can easily revise as I gain more information.

How do I make the revision? By following a broadly inductive policy.

This is another epistemic policy; intuitively, a “higher level” one.

(The relation of one epistemic norm or policy being “of higher level than” another isn’t so much a comparison of the policies themselves as of our relations to them: the manner in which we employ them—or the manner in which we *should* employ them, on the normative realist alternative.)

Claim 1 is that there is a highest level policy that we do, or should, employ.

The conviction that there is such a highest level norm is probably less due to arguments for that position than from an inability to think of a clear alternative to it.

But let’s see if we can construct an argument.

A key premise: that the relation of one norm's being 'of higher level than' (for a particular person) is a partial order: transitive and irreflexive (hence anti-symmetric).

> Seems plausible, but maybe only because the terminology 'higher level than' pre-judges it?

Clarifying the premise: The natural thought is that if I revise my policy about trusting the NY Times, by finding a new domain in which to distrust them,

> (i) this will be because my previous policy of trusting them in this domain came into conflict with other things I believed;
>
> and (ii) I used a general inductive policy of some kind to decide that in this situation it was the policy of trusting the NY Times that should be scrapped.
>
> So the inductive policy overruled the policy about trusting the Times.
>
> **It is the relation of *potentially overruling* (which seems obviously irreflexive) that the "levels" talk takes to be transitive as well.**

Plausibility:

(1) *The weaker anti-symmetry property:* seems hard to imagine that for a normal person, the policy of trust in the NY Times could ever overrule the inductive policy.

> I guess there are people for whom a policy of trust in the Bible overrules their inductive policy, but unless their inductive policy can also overrule their trust in the Bible we have no violation of anti-symmetry.
>
> *Could* an agent regard each of these policies P_{Bib} and P_{Ind} as potentially overruling the other?
>
> > · We can imagine that in some situations the agent takes the Bible as prevailing and in others takes induction as prevailing.
> >
> > · But there would seem to have to be an answer to the question of *when* the first takes precedence and *when* the second does.
> >
> > · And that seems to be a further policy, of higher level than both, that serves for the agent as "the decider" and that can overrule each of them.
> >
> > So we don't count the two lower level policies P_{Bib} and P_{Ind} as potential overrulers of each other, we think of each as potentially overruled only by the decider policy.

(2) Transitivity requires more than anti-symmetry: for instance, it (together with irreflexivity) rules out loops of arbitrary finite size, not just of size two.

> Having no loops is the key: it suffices for *the transitive closure* to be a partial order.

But lack of loops seems plausible on grounds analogous to the ones used for anti-symmetry:

> If in some situations we appear to let policy P_1 overrule P_2, in others let P_2 overrule P_3, ... in others let P_{n-1} overrule P_n, and in others let P_n overrule P_1, then in any situation where the policies come into conflict we have to decide which not to follow.
>
> (I don't mean we consciously decide; we "decide" only in that there are ones that we don't follow in those circumstances.)
>
> What determines which we follow? Seems like a "decider policy", that is of higher level than all of the policies P_1 through P_n.

Let's tentatively grant that this establishes a partial order among a person's policies at a given time, so that we can sensibly speak of one policy as being *of lower level than* another.

Define: P_1 is *of the same or incomparable level as* P_2 to mean: it is neither the case that P_1 is of lower level than P_2 nor that P_2 is of lower level than P_1.

> (There may be principles for deciding in which cases we want to regard P_1 and P_2 as *of the same level*, as opposed to *incomparable*, but the distinction won't matter here.)`

Claim 1 was that at any time, each person possesses a "highest level" epistemic norm. This requires more than just that *potentially overrules* is a partial order.

It rules out, for instance, that for each norm N that the person employs, she possesses a higher norm N* for use in revising N should the situation arise.

But this would require the person to possess infinitely many norms at a given time, which seems prima facie implausible.

The partial order, together with a plausible assumption about our finitude, entails that there must be in each person at each time *maximal norms*: norms for which no other norm in the person at the time is of higher level than it.

Nothing said so far rules out there being more than one maximal norm: if more than one, they are of the same or incomparable levels.

But if there is more than one, why not just regard them not as separate norms, but as aspects of a single norm?

> This seems quite natural *if the maximal norms in question can't come into conflict*.
>
> And it seems plausible that they can't come into conflict: if the norms were to come into conflict, wouldn't we try to resolve the conflict, and wouldn't this require a higher level norm? But by their maximality, such a higher level norm can't exist in this case.

If this is right, we have a whole bunch of maximal norms *that can't come into conflict*, and we may as well regard them as just aspects of a single norm.

An alternative to the idea that the maximum norms can't come into conflict: they can, but we have no norm for resolving the conflict. But here too we could always count these "maximal norms" as just part of a single norm—in this case, one which could lead to inconsistent verdicts under unfavorable conditions.

This completes something of an argument for Claim 1.

"No one ever doubted the existence of God until Samuel Clarke undertook to demonstrate it."

—Anthony Collins.

I think there are some questionable moves in the argument—or rather, one questionable move made at several places in different guises. I'll come back to this, in the final lecture.

But I don't think that the *whole* problem with the incompatible foursome is in Claim 1.

And that's good, because the plausibility of Claim 1 is due less to arguments for it than to the difficulty of finding an adequate description of our epistemic behavior that doesn't countenance a highest level norm.

Claim 2: It isn't possible for a person to rationally revise her highest level epistemic norm under any conditions.

This seems a clear consequence of three premises:

(a) Rational revision requires the use of a norm;

(b) If the rational revision of a norm N went by use of some norm other than N, N couldn't be a "highest level" norm;

and

(c) No norm can dictate its own revision. (Or at least, no norm *that can be used* can do so.)

(a) seems pretty intuitive: if someone revises a norm by some means other than a norm—say a whim, or a blow to the head—then even if the new norm is better than the old, the revision itself is not a rational one.

(b) seems hard to question: it seems to follow from the meaning we've given to "highest level norm".

(c): How could any norms—or any ones that an agent could follow—tell us to revise themselves? Wouldn't following those norms require not following them? This seems incoherent.

On (c):

I'm interpreting (c) to mean that no norm can dictate that it be revised *even in part*.

> It's natural to suppose that if there's rational revision in norms, it's gradual: using a norm P_1, we make a small revision to a norm P_2, then to P_3, and so forth; a large change in norm would take a long chain of such small revisions.
>
> One way to describe a small revision from P_i to P_{i+1} is to say that P_i isn't revising all of itself but only part of itself.
>
> But I'm taking (c) to go against this: P_i dictating a revision of part of itself counts as P_i revising itself, so that if (c) stands, we can't have a chain of small revisions where each P_i dictates that it be replaced by P_{i+1}.

(c) No norm can dictate its own revision [even in part]. (Or at least, no norm *that can be used* can do so.)

The argument for (c) in more detail:

A norm is a policy of some sort—let's say for the moment, it's a precise policy P that dictates what to do under every conceivable situation.

Suppose there were a conceivable situation E such that, were it to arise, P would dictate it's own revision.

That is, P would dictate that from that point on, one should follow some alternative policy P*.

What must be the case for P* to be genuinely an alternative to P, even given the obtaining of E?

> There must be some conceivable situation F that includes E's having obtained, in which P* and P differ in their dictates about what to do: P dictates doing X and P* dictates doing X*.

But in that case, if P dictates following P* it must be inconsistent:

> by hypothesis it dictates doing X,
>
> but it also dictates switching to P* and doing X*.

And presumably an inconsistent norm is one that no one could employ.

That's the initial argument. One way to resist it is in the last step: to say that it's quite possible to employ an inconsistent policy *for a while*, that is, before the inconsistency becomes salient.

Before evaluating this, it's important to be clear what is meant by an inconsistent policy.

A policy that, under certain conditions C, one should simultaneously believe A and believe $\neg A$ is not automatically an inconsistent policy, for two reasons.

> First, it isn't a policy that requires inconsistent belief, it only requires inconsistent belief *should circumstances C arise*.
>
> Second, even a policy that unconditionally requires inconsistent belief can be a perfectly consistent policy, as long as it doesn't simultaneously require not having inconsistent belief.

On this latter:

Having inconsistent beliefs isn't such a terrible thing, even if we recognize the inconsistency: there are ways of working with inconsistency that keeps it from spreading.

(Paraconsistent logics, that say that inconsistencies don't imply everything.)

More important: "Chunk-and-permeate procedures" (Brown and Priest): When one finds oneself with an inconsistent set of beliefs but doesn't know how best to remove the inconsistency, one

(i) breaks the set of beliefs into useful "chunks" ,

(ii) reasons logically within each chunk,

(iii) adopts restrictions about which sort of conclusions drawn within one chunk are allowed to pass to another.

This is a good model of how we steer around recognized inconsistencies within our beliefs. **Following a specific chunk and permeate strategy can be part of a consistent policy for dealing with inconsistent beliefs.**

But if we can work with inconsistent beliefs, can't we work with inconsistent policies? Aren't our policies "written into our heads" in the same way theoretical beliefs are, and doesn't this make them liable to inconsistency too?

I think that is so, for some kinds of policies. But not for any policy that is a candidate for a fundamental norm.

Policies are sets of rules. We can't suppose that all of the rules that an agent employs are explicitly represented in their heads: familiar regress. (We would then need policies or rules for processing the internal representations.)

Some of the policies or rules that we describe a person as employing are merely *implicit in the person's practice*.

But this means that they result from a kind of idealization in the person's practice. And it is at least somewhat natural to suppose that *the process of idealization imposes consistency*.

If so, we can't reasonably suppose that the high-level rules or policies governing a person's epistemic behavior are inconsistent.

And then the idea that we employ an inconsistent policy *for a while*, until the inconsistency becomes salient, is no way around the argument that a policy that dictates its own revision is inconsistent and therefore unemployable.

But in any case, I don't think that allowing for the possibility of inconsistency in fundamental policies is enough to much reduce the plausibility of the case for (c).

Let's grant that we can follow inconsistent fundamental policies. Then we can have an inconsistent policy P that dictates doing X and simultaneously dictates switching to a policy that demands doing something incompatible with X.

> There's no way to both do X and not do X, so the policy in effect leaves us with no idea what to do.
>
> And it's a bit misleading to say that it dictates its own revision: it does, but it equally dictates that it *not* revise itself. (Indeed, on the assumption that it dictates all consequences of what it dictates, then barring a paraconsistent logic it dictates *everything*.)

This doesn't look like a good way to understand the rational revision of logic.

Review:

I made an initial case for Claim 1, the existence (at a given time in a given person) of a "highest level" epistemic norm.

Claim 2, that *if* there is such a highest level norm then it is not rationally revisable, was divided into three sub-claims:

> (a) that rational revision requires the use of a norm,
>
> (b) that a "highest level" norm couldn't be rationally revisable by virtue of its revision being dictated by some other norm,
>
> and (c) that no norm could rationally revise itself.

(c) probably seemed the most contentious of these, and it is the one I mostly discussed. I take it that the discussion has lent some kind of support for Claim 2.

Putting it together with the case for Claim 1, we get the conclusion that at any time a person possesses what we might call an ***ultimate norm***: a highest level norm *that is not rationally revisable*.

****

An initial case only—will revisit in Lecture 6—but I think one with some force.

And common arguments against a highest level norm that isn't rationally revisable strike me as quite weak.

For instance:

A. Any sufficiently powerful norm N must include an inductive policy P_N.

B. For any norm N, one can imagine strong inductive evidence that some other norm N* with a different inductive policy works better than N does.

> (And this is due to the difference in inductive policy; that is, there is no N** that matches the performance of N* while keeping the inductive policy of N).

C. Given such evidence, it would be rational to switch from N to N*.

Why should we believe Claim B, even without its parenthetical clause?

Claim B (truncated): For any norm N, one can imagine strong inductive evidence that some other norm N* with a different inductive policy works better than N does.

Appraisal:

I grant: For any known attempt to formalize our inductive policies in a sufficiently informative way, one can imagine evidence that would undermine the inductive policy as formalized.

> Indeed, for any known attempt to formalize our inductive policies in a sufficiently informative way, *we actually have* evidence that undermines the inductive policy as formalized.

The proper conclusion: all known attempts to formalize our inductive policies fail to do so.

> They may adequately formalize the use of certain kinds of evidence, but they aren't general enough to formalize all the ways that we use empirical evidence.
>
> That's why it's possible to come up with evidence of a kind that they don't take account of, to undermine the dictates of the policy as formalized.

That's hardly a shocker: our actual inductive policies are extremely complicated, and it would be very hard to formalize them in full generality.

There's no reason to extrapolate from the undermining of an *inadequate* formalization of our inductive policies to the undermining of an *adequate* formalization.

Where are we?

I was considering

(I) There is a highest level *inductive policy* that no *empirical considerations* could undermine,

for its bearing on

(II) There is a highest level *norm* that no *rational* considerations could undermine.

Even a strong case in favor of (I) wouldn't obviously provide positive support for (II).

But a case against (I) would seem to extend to a case against (II).

However: The last slide undermines the argument against (I), so as yet we have no case against (II). And I have given a *prima facie* reason in favor of (II).

But (II) implies that *if all of logic is included in the highest level norm, then logic is not rationally revisable*.

However, I'll argue that logic *is* rationally revisable, in Lectures 3 and 4.

One possible reaction is to discriminate between parts of logic: posit a **"core logic"** that is included in our highest level norm, together with various extensions that can be rationally revised.

I won't consider this view in detail, but a serious worry about it is that there is no obvious basis for dividing logic up into "core" and "periphery".

That is, if one is willing to grant that logic is rationally revisable at all, it seems difficult to set off any part of it as immune to rational revision. Consider some examples:

> Considerations of the paradoxes seem to me to support restricting *excluded middle*, *common versions of reductio*, and certain *laws involving the conditional*.
>
> Dialetheists propose that instead of giving up excluded middle we give up *disjunctive syllogism*.
>
> Intuitionists propose giving up not only excluded middle but also the *law of double negation* and one of the *deMorgan laws*.
>
> Vann McGee has famously proposed restricting *modus ponens*, though his example may support giving up *the equivalence of* $A \rightarrow (B \rightarrow C)$ *to* $A \wedge B \rightarrow C$ instead.
>
> Putnam proposed giving up the *distributive law*.
>
> Even a law as obvious seeming as *the inference from* $A \wedge B$ *to* A has sometimes been doubted, e.g. by advocates of dynamic logics in which the second conjunct can undo the commitments of the first.

If our norms allow for serious consideration and perhaps adoption of any one of these under suitable circumstances, what will be left in core logic?

I don't think this entirely settles the matter, for two reasons.

> First, some of these proposals are less serious than others. It isn't clear that we should regard the more frivolous proposals as supporting a case for rational revisability.
>
> Second, there's a different way of conceiving of a "core logic": not as a set of topic neutral logical principles, but as a set of logical principles *together with a domain of application*. The idea: there is no rational revision of those principles *in that domain*.
>
> > For instance, few of the examples above could be used to make a case for restricting classical logic *within arithmetic*. Equivalently: *within proof-theory*. Perhaps it could be argued that all a highest level norm requires is a rationally unrevisable logic for proof theory.

So the issue of a core logic is complicated. I won't in the end find a need to pursue it, because I'll locate the main problems with the inconsistent foursome in Claims 1 and 2.

But there's one other option for those who

> accept Claims 1 and 2,
>
> accept the rational revisability of logic,
>
> and are unwilling to posit a core logic that is not rationally revisable.

That remaining option is to divorce logic completely from our norms.

That option seems to be advocated by Harman. It's the topic of the next lecture.

Preview of Lec. 2

Harman

1. Argues that it's impossible to spell out a believable connection between logic and rational belief, for a battery of reasons

2. Argues for a view of logic that doesn't need one. Viz.:

> logic is the science of what preserves truth by logical necessity.

I'll combat #1 by trying to spell out a believable connection, one that overcomes each of the problems he raises.

Re #2:

I'll argue that the view of logic as the science of what preserves truth by logical necessity can be decisively shown wrong.

And there's no other even initially plausible alternative to the view that there is a special normative role for logic.