How, if at all, can we rationally revise our inductive methods?

Main topics

Two questions (for the initial argument):

1. Are our inductive methods rationally revisable? If so, how?
2. Are they rationally revisable *on inductive grounds*?

Motivation for inductive revisability?

3

One version of the view that they are revisable on inductive grounds:

1. The point of an inductive method is to deliver important truths about what we haven't observed (given other truths as input), without delivering too many falsehoods;
2. We can gather evidence as to how well the methods we employ do this, in comparison to other methods that might be proposed;
3. If we discover that other methods do better than ours, we can try to retrain ourselves to follow those other methods.

This sounds like a rational process.

Does this apply without restriction?

4

This certainly makes sense for sufficiently "low level" inductive methods.

[E.g., methods like "If one discovers that in the recent past most members of species S have had trait P, infer that that will probably be true in the near future too".

That's obviously overturnable, e.g. by evidence that most other related species are undergoing rapid change with respect to related traits.

In overturning it we seem to be employing a broader inductive method, that allows "cross-inductions".]

It might be argued that the model doesn't make sense for *our most general methods*.

E.g., to one's most general methods?

(i) When gathering evidence about how well an inductive method works, one needs to use a more general inductive method.

(ii) So there's no way to inductively revise the most general inductive methods that one employs.

[The more general inductive method required by (i) is absent in this case.]

Of course, (ii) has a presupposition that goes beyond (i): that there is such a thing as *the most general inductive method one employs*.

(Not a most general possible method, just more general than any other *that the agent employs*.)

I have some sympathy for that presupposition, suitably understood, but will defer the issue until later on.

Or even to other sufficiently general ones?

Even without the assumption of a most general method that one employs, there's still a problem about the model of inductive revisability on the earlier slide.

The key claims of that model:

- We can gather evidence as to how well the methods we employ do at delivering important truths without delivering costly falsehoods;
- If we find that the methods we employ do badly (and know of an alternative that would do better), we abandon those methods (in favor of the alternative).

There's an important ambiguity here.

Making the crucial assumption explicit

7

What we can gather evidence for is how well our inductive methods *have done in the past*. (And how well the alternative methods would have done in the past, had we applied them.)

And it is highly plausible that if we think that our current methods will do less well *in the future* than a given alternative, we should switch to the alternative.

Once the italicized clauses are included, it's clear that the model works only if we can reasonably extrapolate from past success or failure to future success or failure.

But this is highly problematic (and not because of inductive skepticism).

"Self-correcting" methods

8

The extrapolation from past failure to future failure looks highly suspect, in the case of any method general enough to provide a reasonable model of a substantial part of our inductive practices.

- For such methods have what we may call (somewhat vaguely) a "self-correcting character": they can be expected to do better as they run longer.
- As they run longer, our background beliefs improve; and with improved background beliefs as inputs, the same method yields better results.

Illustration

Crude illustration: a Carnapian agent whose method is to let his degree of belief in A at a given time be c(A|E), where

- E is all the information it has gathered
- c is an initial credibility function (an "informationless" assignment of conditional degrees of belief).

Illustration continued

Early on, such an agent will presumably do very badly at even simple tasks, because he doesn't have enough evidence.

But as evidence is gathered, future applications of the method become qualitatively different from past applications: e.g. the agent's degrees of belief may become far more in accord with actual frequencies.

[Aside: contrast between 1 & 2 makes talk of "reliability" of such methods hard to interpret.]

One can invent alternative initial credibility functions that would have done better, by predicting all the evidence E; but that doesn't mean the agent would do better in now shifting to one of them.

Clarifying the issue

Obviously some inductive methods are better than others:

- for instance, probably any method that has as of yet been formalized is sensitive only to limited forms of evidence;
- and some methods are sensitive to some forms of evidence in a sub optimal way.

The question is whether it's possible for an agent following a given fairly general method *to gather inductive evidence of* its sub optimality.

I've argued that it isn't straightforward how this would be done: we can't just extrapolate from past success/failure.

And the claim that it can *somehow* be done would require argument.

Interim conclusion

The upshot so far: even without the assumption of a *most* general inductive method that an agent employs, it isn't at all obvious how to revise *fairly* general inductive methods *on inductive grounds*.

The argument can be made stronger if one assumes there is a most general inductive method one employs: then it's hard to see how to inductively revise it.

I'll come back to that. But first:

Non-inductive revision of inductive methods

13

Might we rationally revise our inductive methods by *non*-inductive means?

I'll argue that we could.

A fanciful case: an agent

- whose inductive method M is extremely limited: say, it's a method in Carnap's λ-continuum (or an effective method that approximates to this)
- but who is sufficiently logically sophisticated to develop possible inductive methods that are sensitive to things (e.g. order) which no Carnapian method is sensitive to.

Fanciful example continued

14

The person might then argue that it would be better to employ one of the other methods, M*.

Reason: there are hypotheses about the world which if true, that method could detect while the Carnapian method couldn't.

[Assume that M* isn't so complicated that training oneself to employ it would be impossible.]

In that case, the agent might reasonably train himself to employ the new method. (Rather than continue to simply conditionalize, which is what orthodox Bayesianism dictates.)

Realistic cases? (1)

Are there *realistic* cases of rational revision of inductive methods?

An obstacle to providing one is that the inductive methods we actually employ are immensely complicated, and pretty good.

- It's a huge task to give an accurate description of them, and to investigate their properties.
- It is at least as huge a task to come up with a description of an alternative that has a decent chance of being better, and to investigate its properties.
- As a result, while it's easy to imagine rational revision of our theory about *which inductive method we fundamentally employ*, it's not so easy to imagine rational switching from one inductive method to another.

Realistic cases? (2)

Despite this, I think it may be possible to recognize that our inductive method has a certain feature, and argue that it would be better to alter this feature, and train ourselves to do so.

Examples involving logic

Here's an illustration.

(It may seem like a cheat, and I'd like to have a better one. But I don't, so it will have to do.)

(i) Any inductive practice incorporates a logic.

(ii) The inductive practices most of us follow incorporate classical logic (or some logic relevantly similar to it).

(iii) One can invent alternative logics, and there's a good case to be made that some of these are better.

(iv) To the extent that this is so, it seems we should modify our inductive practice to incorporate such a logic instead of classical.

Inductive methodology and logic

Re (i) ("Any inductive methodology incorporates a logic"): I don't have a general argument, but examples support it.

- Standard Bayesian methodology incorporates classical logic. (Classical theorems get probability 1, classically valid 1 premise inferences preserve probability, etc.)
- You can modify Bayesianism so as to not incorporate classical logic, but different modifications incorporate different logics.
- Any version of hypothetico-deductivism incorporates the logic used in making the deductions.

I'd also argue that you can't understand what it would be to employ a given logic as one's all-purpose logic without understanding how an inductive practice could incorporate it.

Conservativism about logic

19

Many people have opined that one can't rationally change one's deductive logic.

"Not everything can be revised, because something must be used to determine whether a revision is warranted…." Tom Nagel *The Last Word*, p. 65

"There is no intellectual position we can occupy from which it is possible to scrutinize [our logical beliefs] without presupposing them. That is why they are exempt from skepticism. They cannot be put into question by an imaginative process that essentially relies on them." Ibid. p 64

"The only way to revise one's logic is by brain surgery" Louise Anthony (recent lecture)

General approach to revision

20

If this is plausible about logic (the logic incorporated in our inductive methodology), it would seem equally plausible about other aspects of our inductive methodology.

But why is it plausible anywhere?

A general approach to revision: theorize about

1. what possible logics and inductive methods there are
2. which possible ones do best in achieving what we take as the virtues.

In doing 1, we use our current logic; in doing 2, we use our current logic and inductive methods.

If we come to think an alternative logic or inductive method better, we can

3. try to train ourselves to employ it.

Circularity worries

21

Those who think rational revision impossible stress circularity worries in stage 2 (the evaluation of how well methods do).

But the worries are overstated.

In the logic case, it's possible to use one logic L to argue that a theory employing another logic L* is better in achieving truth-oriented goals (such as theoretical simplicity) than any theory employing L. (The argument can usually be given in $L \cap L^*$.)

Admittedly, using L we can also argue directly for the superiority of L over L*: e.g. that L preserves truth and L* doesn't.

But **why must that argument trump the former?**

Indeed, the latter argument seems question-begging (not that there's always anything wrong with that), since L* would rule the opposite.

Immodesty arguments

22

In the case of induction, there are analogs of the question-begging arguments: immodesty arguments, where e.g. each c-function declares itself best. But

1. They depend on non-obvious scoring measures
2. As in the logic case, they don't preclude alternative arguments for the opposite conclusion.

So even if we concede the scoring measure, we have the question: why should the immodesty argument trump arguments for alternative conclusions?

Observational practices

The situation of competing arguments for and against change also arises for observational practices, in "Feyerabend cases":

- cases where the old observational practice is laden with a theory that can be questioned.

Here too one *could* use the old observational practice to dismiss the new theory.

"We've observed thousands of bodies falling in straight lines, whereas Copernican theory says they fall in curved arcs because of the spinning earth; so Copernican theory is wrong."

No trumping

Arguments of this form, though question begging, don't always completely lack force. [Other cases, incl. brain in vat.] I don't say we should discount them.

But the proper response to them involves ensuring that they don't trump arguments for competing conclusions.

Similarly in the case of logic and inductive methodology.

Dealing with inconsistencies

25

One can say something vague but useful about how the competing arguments are to be handled.

("Develop each of the alternatives, even if they conflict with prior observational practice, logic, or inductive practices. See which does best.")

But if one wants to develop something more precise in these cases of conceptual change, I think one needs to take seriously that during conceptual change, different aspects of one's theoretical state are pushing in incompatible directions.

So it would be good to bring in models of how to deal with inconsistent information.

Chunk and permeate

26

A model of how we do so was suggested by Graham Priest and Bryson Brown, in their paper "Chunk and Permeate". The general idea:

- Our cognitive processes are divided into "chunks", within which we reason using our logic (in the present context, the old logic).
- Instead of allowing free passage of information between chunks, we impose restrictions.

Illustration: reasoning with infinitesimals.

This model probably needs to be generalized a bit if it is to be applied to revision of fundamental practices (e.g. logic, or observational practices).

Generalizing chunk & permeate

27

For this, we probably want an indeterministic model of mentality, where something like chance plays a role both in

- what theories (e.g. logical theories) one thinks up, and
- what evaluation one comes to of the respective merits of the theories.

Presumably focusing on one chunk and becoming influenced by its conclusions will diminish the influence of incompatible chunks.

So we don't want the *fixed* restrictions on information transfer assumed by Priest and Brown.

But whatever the details, there is little doubt that we have rational ways of dealing with inconsistent premises.

Application to avoid trumping

28

And we can then use this way of dealing with inconsistency, in the case where

- we have good arguments for substantial change of logic, inductive methodology, or observational practices
- competing with the obvious ("question begging") arguments against such change.

Are psychological models beside the point?

29

Objection: a mental model of how we deal with inconsistency would tell us how we *do* change logic, not how we *should*. It would leave open the question of whether acting in this way is *rational*.

Reply: Once we have a model of how we *do* change logic (or even of how we *might*), the question of whether the model makes the change rational is simply a question of evaluation:

Is acting in accordance with the model a good thing (or would it be, if we don't actually act that way).

To answer this, we must compare it to alternative models. And it's hard to believe that the dogmatic models could win.

Summary of Overall Argument

30

Reprise: The key features of rational revision of inductive methodology are:

1. coming up with an alternative methodology (described in sufficient detail)
2. arguing for merits of new methodology over the old (using the old, or what's common between the two)
3. retraining ourselves.

Stage 2 is complicated: the old methodology will always have (at least cheap) arguments that it is superior to the new.

But this doesn't prevent arguments in the other direction. We need to weigh the arguments on each side.

We have intuitive ways of doing this, but a formal theory of how we deal with inconsistent information would be nice.

Revision on *inductive* grounds looks dubious.

Back to "one's most general" methods

31

I've given the overall argument, but would now like to come back to the question of whether there is a most general inductive method that an agent employs [a method *fundamental for the agent*].

- Seems of interest in its own right
- A positive answer would seem to support my conclusion that you can't *inductively* revise general inductive rules.

But there's a complication: won't arguments that there is a most general *inductive method* show equally that there are most general *rational rules* that an agent employs [rational rules fundamental to the agent?]

And mightn't this support the view that inductive rules can't be *rationally* revised? If so, I'm in trouble.

Rational revision of rational rules?

32

The last is too quick: if an agent employs fundamental inductive rules and also fundamental rational rules (and rationality includes induction), there seem to be three questions:

1. Can the fundamental inductive rules be inductively revised?
2. Can the fundamental inductive rules be rationally revised?
3. Can the fundamental rational rules be rationally revised?

There is no obvious tension between a 'No' to 1 and a 'Yes' to 2. And it's 2, not 3, that I've maintained.

Still, there seems to be tension between a 'No' to 1 and a 'Yes' to 3. And there may be some discomfort in a 'Yes' to 3.

Clarification

33

Of course, the question about revision of fundamental inductive or rational methods (on the assumption that there are such methods) isn't

- whether an agent could inductively or rationally revise *her opinion about what those methods are*.

Rather, the issue is

- whether an agent could switch from actually employing one fundamental method to actually employing a different one, on inductive or other rational grounds.

The "no inductive grounds for switch" argument

34

The argument that an agent could have no rational grounds for switching fundamental methods:

- If she employs a fundamental method, she shouldn't revise any method unless her fundamental method tells her to.
- Unless the fundamental method is inconsistent, it can't tell her to stop following itself and to follow some other method instead.

(Focusing on induction:) The method M might indeed say:

- Given evidence E, then *unless you also have evidence F*, believe B; and if you do have evidence F in addition to E, determine what to believe by alternative method M*.

But that isn't M dictating its own revision, it's simply a case of M telling us to treat F as a defeater to B.

The argument elaborated

35

A method M *would* be dictating its own revision if it said:

- Given evidence E, believe B *whatever the circumstances*.
- But if you have evidence F as well as E, determine what to believe by method M*

  [where M* might deliver non-belief in B when applied to evidence E and F].

But obviously such an M is inconsistent.

The argument's upshot

36

The upshot:

- It seems hard to imagine how an inductive method can dictate its own revision.
- So if there is an inductive method that is *the most general one we employ*, then we can't rationally revise it *by any inductive method we employ*.

(A being that used more powerful inductive methods than any we employ might argue, using *those* methods, that *ours* shouldn't be accepted.

But she wouldn't be able to inductively revise her own most general methods.)

And the same would seem to go for rational revision of rules of rationality.

Clarifications?

37

It would be good to clarify what might be meant by an inductive method, and rules of rationality more generally.

Also, what's meant by an agent's "most general" inductive method or rules of rationality.

But I don't know how to be as clear as I'd like.

"Inductive method": clarifications (1)

38

Basically, an inductive method directs what you are to believe, and to what degree, given your observations. But

- It needn't be deterministic: at certain points it can allow different choices of what to believe. (It needn't even supply determinate transition probabilities.)
- Its dictates needn't be determined solely by the evidence and our other beliefs: e.g., it could depend on the class of hypotheses that have been proposed or that the agent has thought of.
- There are also issues about how "the evidence" is described, given Feyerabend's point that revisions of one's beliefs, or even of what one takes the live possibilities to be, can lead one to reject previous formulations of the evidence.

"Inductive method": clarifications (2)

There are issues at the output end too: e.g.

- Should degree of belief be viewed in something like Bayesian terms, or in terms of "spheres of entrenchment", or in some other way?
- Should we include in the outputs additional things, such as decisions about what further evidence to gather?

For more general rules of rationality, there are many more such questions.

Computational feasibility

Presumably the methods must be *computationally feasible*.

One often takes Bayesian probability theory with a given initial probability distribution (or class of such) as a toy example. (And it is a useful one since there are many respects in which it gives qualitatively correct behavior.)

But taken literally, it offends badly against computational feasibility, since reasonable probability functions are never even Turing computable. (By a generalization of Church's Theorem.)

[Viz: no recursive set includes all theorems of Robinson arithmetic while excluding all anti-theorems.]

Indefinite extensibility? (1)

41

Some might try to argue that the requirement of computational effectiveness precludes there being a most general method that an agent employs.

The claim might be made by analogy to, or as a consequence of, a claim often made about Godel's theorem (e.g. Dummett):

- that it shows that our abilities to ascertain arithmetic truths (our "concept of natural number") cannot be given by a single formal system, but is instead to be given by a formal system *plus a rule for expanding it*.

But (as I think is generally recognized?) this is highly dubious: it depends on assuming that if a system codifies our mathematical abilities, we must have the ability to recognize that system as sound.

Indefinite extensibility? (2)

42

Moreover, the view itself is of dubious coherence:

- If our abilities are given by a formal system plus a rule for expanding it, mustn't that rule for expanding it be recursive?
- And if so, isn't the view equivalent to a view that involves a single formal system, viz. the resultant of reiterating the expansion for as long as one can?

[Any effective process of transfinite iteration of the rule must have a recursive ordinal as a limit. Transcending the process requires the use of a more complicated ordinal, which will typically require more conceptual sophistication.]

Indefinite extensibility for induction (1)

This is a worry in the inductive case too:

Consider any model that purports to posit an alternative to there being a most general inductive method.

- E.g., a sequence of more and more general ones, with an effective rule for going from any one to the next when circumstances warrant.

The worry is that any precise version of this will be equivalent to a single most general method.

(Roughly, the union of the methods given in the alternative description.)

Indefinite extensibility for induction (2)

For instance, a natural attempt to avoid the conclusion that we employ an inductively unrevisable method is:

- we employ several distinct inductive methods M_i, where any one could be inductively revised by employing one of the others.

(They needn't fall in a hierarchy of increasing generality.)

But this leads to the idea that the different M_i can lead to different conclusions about a given case: if one follows M_1 one gets one verdict, if one follows M_2 one gets the verdict that one shouldn't follow M_1 but do something else.

And this raises the question of whether we employ a tie breaking procedure.

Tie-breakers?

45

If we do employ a tie-breaking procedure, shouldn't a proper description of our method be:

- Employ the union of the M_i *except when they conflict*, and in that case *use the tie-breaking procedure*?

And if we don't employ a tie-breaking procedure, but are free to settle conflicts as we please, shouldn't a proper description of our method be:

- Employ the union of the M_i *except when they conflict*, and in that case *follow whichever we like*?

Conclusion on "our most general method"

46

I tentatively conclude:

- We don't have a clear alternative to the idea of an agent's most general inductive method, or of an agent's most general rational method.

I think this does tend to reinforce the claim against inductive revisability of some inductive rules: these most general rules are not inductively revisable.

At the same time, I think the conclusion does *not* support the view that there are most general rules of *rationality* that are not *rationally* revisable.

Why?

Qualification: indeterminacy as to what that method is?

47

First I need to clarify the conclusion in the inductive case.

Regarding the conclusion of an agent's most general inductive method, bear in mind:

Talk of a person's inductive method is an idealization, from a person's inductive practices.

The practices are messy, and it is by no means obvious that there is a uniquely best idealization.

Similarly for rational rules.

Indeed, *at least in the latter case*, there's a strong reason for thinking that there's no uniquely best idealization: we can idealize "on different levels".

A reason for believing it indeterminate

48

That is, we can make different decisions as to which computational limitations to take into account in our idealization.

From the viewpoint that takes account of certain computational limitations, a method that doesn't will seem deficient because unusable.

From the viewpoint that doesn't, a method that does will seem to license irrational behavior due to imperfections in the hardware.

We don't have a clear enough notion of rationality to choose between these.

(Not the same as the view considered before, according to which we have an indefinitely extensible method not describable by a single formal system.)

Conflicts between "levels"

49

It's conceivable that for any clearly reasonable level of idealization:

a person's most general rational method at that level of idealization will be rationally revisable by the person's method at some *fairly* reasonable lower level of idealization.

The possibility turns on the vagueness of 'reasonable idealization'.

False dichotomy

50

But if we precisify 'reasonable', the earlier argument **seems** to apply:

In a case of conflict between a lower level of idealization and a higher:

1. If we base our rationality judgement on the dictates of the lower, then the "higher level method" is at best a default method.
2. If we base our rationality judgements on the higher, then the "lower level method" brings in factors (e.g. computational limitations) unrelated to rationality.

However: I think this is a false dichotomy; our concept of rationality doesn't work this way.

The role of rules

In giving a model of a rational agent, it's natural (inevitable??) to state the model in terms of rules that the agent follows.

[Not explicitly formulated rules, because of the familiar regress.]

As above, an actual agent will be described by different such models at different levels of idealization.

But we should view these as psychological models of a rational agent, rather than thinking that they give explications of our term 'rational'.

Example

Crude example: a 2 tape Turing machine.

- The top tape contains instructions for inductive or rational method M. It is appealed to constantly in rewriting the bottom tape used for ordinary practice.
- Only in very exceptional circumstances do the overall instructions of the machine dictate rewriting the top tape.

From the point of view of the overall instructions (the machine table), M is merely a default program that has defeaters.

But the machine table is far too "low level" to be naturally viewed as describing "rules of rationality". The only natural candidate for that (before any rewrite) is method M.

Rationality as evaluative, not descriptive

53

If the overall instructions do lead the machine to rewrite method M, due to M "getting into trouble", should the change can't as "rational"?

If we insist on using M to judge rationality, it isn't.

But we shouldn't: in calling something rational, we're evaluating it. The etiology of the rewrite is *relevant* to the evaluation (a rewrite with a good result could be irrational because of how it came about), but there's no reason to think that there's an etiological reduction of the rationality claim.

Application to the 2-tape TM

54

In the case I've described, the change came about by normal operation of the machine, not a malfunction.

To the extent we deem the change a good one, we'll deem it rational.

Doing this does not require taking the machine table to give the agent's "rules of rationality". (They might in other imaginable circumstances lead to changes we'd regard as unfortunate, and 'irrational".)

For the rationality claim is our evaluation, it doesn't describe how the agent made the revision.

Fundamental rules of rationality?

55

I take it that this example has realistic variants. Should we describe cases like it as

- ones where there are fundamental rules of rationality that can be rationally revised?
- or as ones where there are no fundamental rules of rationality?

Seems like a verbal issue.

****

I doubt that the situation with 'rational' applies to 'inductively supported': the latter seems to me a more rule bound notion. Topic for discussion period?

Tentative Overall Conclusions

56

In summary, I'm inclined to the following:

1. It probably makes sense to speak of an agent's fundamental inductive rules (though it may be indeterminate what they are).
2. It probably doesn't make sense to say that they can be inductively revised, though they can be rationally revised.
3. Indeed, it's hard to give a model on which "self-correcting" rules can be inductively revised.
4. Whether to speak of "fundamental rules of rationality" seems a verbal question (though if one does, one should recognize indeterminacy as to what they are).
5. But to the extent that there are plausible candidates, it is doubtful that they are immune from rational revision.