

Induction and Supposition

Brian Weatherson

Here's a fairly quick argument that there is contingent a priori knowledge. Assume there are some ampliative inference rules. Since the alternative appears to be inductive scepticism, this seems like a safe enough assumption. Such a rule will, since it is ampliative, licence some particular inference *From A infer B* where *A* does not entail *B*. That's just what it is for the rule to be ampliative. Now run that rule inside suppositional reasoning. In particular, first assume *A*, then via this rule infer *B*. Now do a step of \rightarrow -introduction, inferring $A \rightarrow B$ and discharging the assumption *A*. Since *A* does not entail *B*, this will be contingent, and since it rests on a sound inference with no (undischarged) assumptions, it is a priori knowledge.

This argument is hardly new. It is part of the argument in some recent papers promoting contingent a priori knowledge, such as ? and ?. But it is an intriguingly quick argument for a stunning philosophical conclusion, one that seems to rely on few dubious steps. I'm going to argue that it fails for a quite interesting reason. At least in natural deduction systems, some inferential rules (such as \forall -introduction) have restrictions on when they can be applied. I'm going to argue that ampliative reasoning rules cannot, in general, be applied inside the scope of suppositions, and that is why the above argument fails.

I'll argue for this conclusion by showing that a very weak ampliative rule leads, when combined with some other plausible principles, to absurd conclusions if it is applied inside the scope of suppositions. If even a weak ampliative rule cannot be used suppositionally, then it plausibly follows that no ampliative rule can be used suppositionally. The construction I'm going to use to show this is quite similar to one used by Sinan Dogramaci in his ?, though as we'll see at the end Dogramaci and I have different views about what to take away from these arguments.

Some people might think we have already seen an argument that ampliative inference rules fail in suppositional reasoning. If these rules are allowed, then we have contingent a priori knowledge, and this is implausible. I don't believe this argument works, since I think there are *other* arguments for contingent a priori knowledge. (Some of them are in the above cited papers.) So it is a live question whether this quick argument for contingent a priori knowledge works.

Here's the main argument. If any ampliative inference is justified, I think the following rule, called 'R99', is justified, since this is a very weak form of an inductive inference.

R99 From *Over 99% of Xs are Ys* and *a is X* infer *a is Y* unless there is some *Z* such that it is provable from the undischarged assumptions that *a is X and Z* and *Less than 99% of things that are both Xs and Zs are Ys*.

† In progress.

Note that the rule does not say that 99% of *observed* Xs are Ys , but that 99% of *all* Xs are Ys . So this seems like a very plausible inference; it really is just making an inference within the distribution, not outside it. And it is explicitly qualified to deal with defeaters. And yet even this rule, when applied inside the scope of suppositions, can lead to disaster.

In the following proof, we'll write ' $99(F, G)$ ' for *Over 99% of Fs are Gs* to shorten the presentation. And to make the rule a little weaker, we'll say that $99(F, G)$ is *false* when there are no Fs . We'll write FH for the predicate taken by conjoining F and H . So $99(FH, \neg G)$ means, in English, that over 99% of things that are both F and H are $\neg G$, and that at least one thing is both F and H . Finally, we'll use second-order variables X, Y, Z with the obvious introduction and elimination rules and let I be the predicate *is self-identical*.

$99(F, G) \wedge Fa$	assumption	(1.1)
$99(F, G)$	\wedge -elimination, (1)	(1.2)
Fa	\wedge -elimination, (1)	(1.3)
Ga	$R99, (2), (3)$	(1.4)
$(99(F, G) \wedge Fa) \rightarrow Ga$	(1) – (4), discharging (1)	(1.5)
$99(FH, \neg G) \wedge (Fa \wedge Ha)$	assumption	(1.6)
$99(FH, \neg G)$	\wedge -elimination, (6)	(1.7)
$Fa \wedge Ha$	\wedge -elimination, (6)	(1.8)
$\neg Ga$	$R99, (7), (8)$	(1.9)
$(99(FH, \neg G) \wedge (Fa \wedge Ha)) \rightarrow \neg Ga$	(6) – (9), discharging (6)	(1.10)
$\neg((99(FH, \neg G) \wedge (Fa \wedge Ha)) \wedge (99(F, G) \wedge Fa))$	classical logic, (5), (10)	(1.11)
$\neg((99(FH, \neg G) \wedge 99(F, G)) \wedge (Fa \wedge Ha))$	simplifying (11)	(1.12)
$(99(FH, \neg G) \wedge 99(F, G)) \rightarrow \neg(Fa \wedge Ha)$	classical logic, (12)	(1.13)
$(99(FH, \neg G) \wedge 99(F, G)) \rightarrow \forall x \neg(Fx \wedge Hx)$	\forall -introduction, (13)	(1.14)
$99(FH, \neg G) \rightarrow \exists x(Fx \wedge Hx)$	definition '99'	(1.15)
$(99(FH, \neg G) \wedge 99(F, G)) \rightarrow \exists x(Fx \wedge Hx)$	classical logic, (15)	(1.16)
$\neg(99(FH, \neg G) \wedge 99(F, G))$	classical logic, (14), (16)	(1.17)
$99(F, G) \rightarrow \neg 99(FH, \neg G)$	classical logic, (17)	(1.18)

This is already a bad result, but worse is to come. Since F, G, H are arbitrary we can infer...

$$\forall X, Y, Z(99(X, Y) \rightarrow \neg 99(XZ, \neg Y)) \quad \text{second order } \forall\text{-introduction, (18)} \\ (1.19)$$

Now substitute I for X and $\neg Y$ for Z .

$$\forall Y(99(I, Y) \rightarrow \neg 99(\neg Y, \neg Y)) \quad \text{second order } \forall\text{-elimination, (19)} \\ (1.20)$$

And that can only be true in worlds with less than 101 individuals. For if there were 101 (or more) individuals, there would be a predicate Y that applied to all but one of them, and then both $99(I, Y)$ and $99(\neg Y, \neg Y)$ would be true.

Note that the only assumptions made in the proof are discharged. So we have an a priori reason to believe whatever conclusions are drawn. As we saw in the first paragraph, these may be contingent conclusions, but this is supposed to be a method for deriving contingent a priori conclusions. So if we can use R99 in the scope of suppositional reasoning, and every other rule in the argument is correctly used, then we have an a priori reason to believe that there are fewer than 101 individuals in the world. Obviously this couldn't be a priori *knowledge* since it isn't *true*, but it is an a priori justified belief.

And that's obviously crazy. We don't have any reason, a priori, to believe the world has that few things in it. Maybe some diehard Occamists will think that a priori we should think the world has 1 thing in it, or 0 if that's possible, but I suspect most will agree that this is a misuse of Occam's Razor. Really, if the proof goes through, we have an argument that R99 cannot be used inside the scope of a supposition.

And it seems the proof does work. At the time we use R99, the *only* live assumption is that a from a particular group, and something about the distribution of G ness in that group. The assumptions couldn't prove anything stronger. The only remotely controversial step after that is the \forall -introduction in step 14. But since there are no undischarged assumptions involving a , indeed there are no undischarged assumptions at all, at this step, it is hard to see why this would fail. For a more positive reason, note that we could replicate the proof for any other name. That is, by repeating steps (1)-(13), we could easily prove $(99(FH, \neg G) \wedge 99(F, G)) \rightarrow \neg(Fb \wedge Hb)$ or $(99(FH, \neg G) \wedge 99(F, G)) \rightarrow \neg(Fc \wedge Hc)$, or anything else we wanted to prove. That's the usual defence of \forall -introduction, so we should be able to infer the universal here.

(Let me set aside one small point. I've gone from $p \rightarrow \phi a$ to $p \rightarrow \forall x \phi x$ rather than $\forall x(p \rightarrow \phi x)$ simply because (a) it's easier to interpret, and (b) it shortens the proof. But the latter two sentences here are classically equivalent, so this shouldn't make a difference to the cogency of the proof.)

Perhaps it will be objected that line (14) is a mistake because although we can prove every instance of the universal quantifier, inferring the universal version creates an undue aggregation of risks.¹ Thinking about this probabilistically, even if line (13) is very probable, and it would still be probable if a were replaced with b , c or any other name, it doesn't follow that the universal is very probable. But I think

¹? blames the \forall -introduction step in the version of the similar proof that he uses. One of his objections is the probabilistic objection I'm making here. Another is, I think, a general sense that rules we learn in logic class, like \forall -introduction, are less plausible than intuitive modes of reasoning like statistical inference inside a conditional proof. I discuss both of these objections in turn.

this is to confuse defeasible reasoning with probabilistic reasoning. The only way to implement this restriction on making inferences that aggregate risk would be to prevent us making any inference where the conclusion was less probable than the premises. That will rule out uses of \forall -introduction as at (14). But it will also rule out \wedge -introduction, and indeed any other inference with more than 1 input step. If we're worried about risk aggregation, we shouldn't even allow the inference to (11), which could be less probable than each of the steps that preceded it. To impose such a restriction would be to cripple natural deduction.

A determined Bayesian might agree at this point. Such a Bayesian will say that the problem here is that we haven't reasoned probabilistically all along. Perhaps that's right. But if that's so, then R99, and all other ampliative rules like it, must be scrapped. Really the Bayesian should want us to infer from *Over 99% of Xs are Ys* and *a is X* that the probability that *a is Y* is over 0.99. And that doesn't lead to disaster. But that just is to deny the existence of ampliative inference rules. What we've ended up with by following through this objection is a much more negative position than the one taken in this paper. I've argued that ampliative inference rules can't be applied inside the scope of suppositions. The Bayesian, or at least the Bayesian who objects to the use of \forall -introduction at (14), thinks they can't be used *anywhere*. Now such a Bayesian may go on to say why this isn't incompatible with inductive knowledge, or they might entreat us to do away with traditional notions like 'knowledge' and replace them with notions like 'high probability'. Either way, it isn't a threat to the position argued here.

We can make the argument of the last two paragraphs clearer by considering five distinct positions about ampliative inference.

1. There is no cogent ampliative inference, and hence all knowledge is deductive consequences of facts of which we are 'directly' aware. Depending on how liberal we are with this notion of directness, this kind of position will allow quite a bit of knowledge gained through perception, testimony, memory and other sources, but it does not allow non-trivial knowledge about the future.
2. There is no cogent ampliative inference, but we can gain knowledge about the future. That's because we know (not via prior ampliative inference) various conditionals of the form *If the past is this way then the future will be that way*, and via such conditionals and non-ampliative inferential rules we can deduce facts about the future.
3. There is cogent ampliative inference, but it is not rule governed the way non-ampliative inference appears to be. This position is a kind of particularism about ampliative inference.
4. There is cogent and rule-governed ampliative inference, but ampliative inferential rules behave differently inside and outside the scope of suppositions. In this respect, the rules are like \forall -introduction and necessitation, which have constraints on when they can be applied, and unlike, say, \wedge -elimination.
5. There is cogent and rule-governed ampliative inference, and ampliative inferential rules do not behave differently inside and outside the scope of suppositions.

In particular, we can use ampliative inferential rules inside the scope of suppositions in order to generate contingent a priori knowledge of conditionals.

My aim here has been to argue against option 5. I take option 1 to be highly implausible, though it isn't entirely without adherents. The overall tenor of my remarks has been to push towards option 4, but I haven't said anything against options 2 and 3. Now if we try to fit the Bayesian into this framework, I think it is clear that they have a version of option 2. Updating by conditionalisation is just the probabilistic equivalent of updating by \rightarrow -elimination; both the person who believes in option 2 and the devotee of conditionalisation thinks the conditional structure of our thought is epistemologically prior to empirical evidence, and the role of evidence is to move us within this structure. I have deep doubts about this position, but those doubts are irrelevant to this paper. The point here is to argue against option 5. And a probabilistic, or Bayesian, objection to my argument isn't really of any help, because once we take the Bayesian position on board we end up with a more radical objection to option 5 to mine, i.e. we end up with option 2.

It might be argued though that this defence of line (14) is too theoretical. The problem is not that (14) makes some particular probabilistic error. Rather, the problem is that the conclusion is absurd, and one of the rules must be false. Since the steps of \forall -introduction at lines (14) and (19) are the least plausible steps intuitively, we should locate the error there. This is an important objection, but I think it is a misdiagnosis of the problem.

For one thing, dropping \forall -introduction in these cases is very odd as well. It's quite counterintuitive to say that for any given object o we can derive that it satisfies $\lambda x.(99(FH, \neg G) \wedge 99(F, G)) \rightarrow x \neg(Fx \wedge Hx)$, and we can know this, but we can't go on to infer $\forall x((99(FH, \neg G) \wedge 99(F, G)) \rightarrow \neg(Fx \wedge Hx))$. We might ask what we're waiting for? For another thing, this seems to locate the mistake *too late* in the proof. It's very odd that we can infer that for any predicates F, G, H and any object o , that o satisfies $\lambda x.(99(FH, \neg G) \wedge 99(F, G)) \rightarrow \neg(Fx \wedge Hx)$. Even if we are barred for some reason for collecting these judgments into a universal claim, the fact that we can make each of them seems already too quick.²

And, at risk of blatantly begging questions, it seems to me we should be very suspicious of the quick argument for contingent a priori justification from the first paragraph. This is not to say that there aren't *any* arguments for the contingent a priori. Perhaps reflections on the nature of natural kinds, and our relation to them, could motivate the contingent a priori a la ?. Or perhaps reflections we see in papers such as ? or ? could drive us to think to avoid external world scepticism we need some kind of contingent a priori. But the contingent a priori has always been controversial in philosophy, so a view that makes one side of the controversy obviously

²The oddity of these conclusions is why the big proof early in the paper goes via H rather than inferring $99(F, G) \rightarrow \forall x(Fx \rightarrow Gx)$ directly from line (5), and then, via a second-order universal introduction, inferring that there must be at most 101 objects. Many people will think that if there is a problem here, it is in one of the steps of universal quantification. Indeed, some people I've spoken to think the intuitive problem with that argument is merely the second-order generalisation, since $99(F, G) \rightarrow \forall x(Fx \rightarrow Gx)$ is a priori justified. I think that's wrong, but I wanted an argument where the first line before a universal introduction was more clearly unintuitive.

correct is counterintuitive. Since the ability to use rules like R99 inside the scope of suppositional reasoning would make one side of the debate trivially correct—line (5) is already an example of the contingent a priori—that suggests the rule is counterintuitive.

An alternative objection is that R99 is too strong, because it doesn't restrict its scope to projectable predicates. It isn't immediately obvious such a restriction is needed. After all, R99 is really a form of *injection* not *projection*, since we are inferring from things we know about the class to things we don't antecedently know about the individual. But perhaps this argument shows that, despite this fact, we need a projectability-like restriction on statistical rules like R99.

It would be too restrictive to say that R99 applies only when F and G are projectable. Consider a case where F and $\neg G$ are projectable, but G is not.³ And assume we know that $99(F, G)$. Then we know that less than 1% of F s are $\neg G$ s. And, since F and $\neg G$ are projectable, that presumably means we have good reason to believe that the next F will not be $\neg G$, i.e., it will be a G . So if F and $\neg G$ are projectable, then R99 looks like a good rule.

And this is enough to lead to disaster. Let F be *is an animal*. Let G be any predicate of the form *is not an S*, where S is a species, and let H be S . I assume that *is an animal* and S are projectable; in any case, they are predicates that we project with all the time. Then the kind of reasoning above lets us get to line (18), which says (after substitutions) $99(F, \neg S) \rightarrow \neg 99(FS, S)$. But $\neg 99(FS, S)$ is true only if there are no FS s, i.e., there are no S s. So there can't be an extant species such that more than 99% of animals are from other species. And from that it follows immediately that there can be at most 100 species. But it is absurd to have an a priori argument that there are at most 100 species of animal in the world. So even a restricted version of R99, one that is sensitive to projectability considerations, still yields an absurd result. I claim the absurdity is from applying R99 inside suppositional reasoning.

Finally, some discussants have argued that it is counterintuitive if we can't, in everyday situations, know claims of the form $A \rightarrow B$, where in our actual situation A would be outstanding, if non-conclusive, evidence for B . But the fact that we can't use ampliative rules in suppositional reasoning hardly entails that conclusion. For one thing, often A is outstanding evidence for B because we antecedently know $A \rightarrow B$. For another, we can sometimes *deduce* that it is rational to believe $A \rightarrow B$. If it is generally acceptable to infer $A \rightarrow B$ from the rationality of believing $A \rightarrow B$, and I think it is, then we can derive $A \rightarrow B$ while only using an ampliative rule in a non-suppositional context. That is the form of argument that's at the centre of the reasoning in ? , and it isn't threatened here. So the restriction I'm suggesting doesn't yield any kind of invidious scepticism.

The upshot of these reflections is that there is no plausible position which holds that rules like R99 can be applied inside the scope of a supposition. Either the argument here shows that such a use of R99 leads to absurdity, or it is a mistake to think of rules like R99 as rules of inference, rather than shorthands for probabilistic

³In general I think the negations of projectable predicates are not projectable.

rules. And if that's right, then the quick argument for contingent a priori knowledge discussed in the first paragraph can't succeed.

April 14, 2010

References

- Dogramaci, Sinan. 2010. "Knowledge of Validity." *Noûs* Forthcoming.
- Hawthorne, John. 2002. "Deeply Contingent A Priori Knowledge." *Philosophy and Phenomenological Research* 65:247–269.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.
- Weatherson, Brian. 2005. "Scepticism, Rationalism and Externalism." *Oxford Studies in Epistemology* 1:311–331.
- White, Roger. 2006. "Problems for Dogmatism." *Philosophical Studies* 131:525–557.
- Wright, Crispin. 2004. "Warrant for Nothing (And Foundations for Free)?" *Proceedings of the Aristotelian Society, Supplementary Volume* 78:167–212.