doi: 10.1111/rati.12129

INDUCTION, NORMALITY AND REASONING WITH ARBITRARY OBJECTS

Markos Valaris

Abstract

This paper concerns the apparent fact — discussed by Sinan Dogramaci (2010) and Brian Weatherson (2012) — that inductive reasoning often interacts in disastrous ways with patterns of reasoning that seem perfectly fine in the deductive case. In contrast to Dogramaci's and Weatherson's own suggestions, I argue that these cases show that we cannot reason inductively about arbitrary objects. Moreover, as I argue, this prohibition is neatly explained by a certain hypothesis about the rational basis of inductive reasoning — namely, the hypothesis that inductive reasoning is fundamentally reasoning about what normally happens (in a non-statistical sense).

1. Introduction

We often reason from things that we believe. On the other hand, we often also reason from things that we do not believe but merely *suppose* — as when, for example, you suppose something in order to prove its negation by *reductio ad absurdum*, or when you reason hypothetically to a conditional conclusion. Moreover, it is tempting to think that the norms of good reasoning are much the same in both cases: intuitively, once you have adopted a supposition you are then rationally permitted to use it much as you would your beliefs.[1] Tempting though this thought might be, Brian Weatherson (2012) has recently argued that it is false: on Weatherson's view, inductive reasoning within the scope of suppositions is rationally prohibited.

Weatherson arrives at this conclusion in response to Sinan Dogramaci's (2010) observation that inductive reasoning interacts disastrously with reasoning patterns that seem fine in the deductive case. While Dogramaci suggests that the problem has

[1] There seem to be propositions that you can coherently *suppose* but not coherently *believe*. Propositions expressed by Moorean sentences such as '*P* but I don't believe that *P*' are examples. The intuitive idea sketched in the text would need to be qualified to deal with such cases. The issues discussed in this paper are independent of such complications.

to do with universal generalisation ('UG'), Weatherson argues that it is much more general. As I explain in Section 2 of this paper, I agree with Weatherson that Dogramaci is wrong to focus on UG in particular as the source of the trouble; at the same time, however, Weatherson's own blanket ban on inductive reasoning within suppositions is unmotivated. I argue for a more moderate claim: inductive reasoning *about arbitrary objects* is rationally prohibited.

As we shall see, this restriction is naturally explained by an independently attractive conception of the rational basis of inductive reasoning. In particular, it flows naturally from taking inductive reasoning to be reasoning about *what normally happens.* Section 3 of this paper is devoted to motivating and explaining this idea, and to showing how it addresses our puzzle.

2. Dogramaci's Puzzle

Let us begin by comparing a couple of cases

> *Case 1.* Seeing his friend Raji walk out of the examination room with a big smile on her face, Bob infers that she did well on her exam.

On a natural reading of this story, Bob's inference is not deductively valid: it seems plausible that the hypothesis that Raji looks happy despite having done poorly on her exam is consistent with everything that Bob knows prior to his inference.[2] Nevertheless, Bob's inference seems rational: other things being equal, it looks like a source of justified belief (and possibly even knowledge). For contrast, suppose that Tom reasons as follows:

> *Case 2.* Tom begins by introducing an arbitrary person, say *a*. He then makes the supposition that *a* just took an exam, and now looks happy. Then he reasons that *a* did well on the exam. Discharging the supposition, he infers that if *a* just took an exam and now looks happy, then *a* did well on the exam. Noting that

[2] This can be challenged: Bird (2005) for instance argues that knowledge-generating abductive inferences are a species of deduction. I cannot engage with this concern at any length here; but even if we accept that you cannot *know P* by inference unless what you antecedently know is inconsistent with the falsity of *P*, it seems intuitively clear that sometimes it can be *rational* for you to infer *P* even if this is not so. This is enough to generate the puzzles.

a is an arbitrary person, he generalises: *everyone* who looks happy after an exam has done well on the exam.

Intuitively, this piece of reasoning is epistemically irrational: no one would be justified in believing that everyone who looks happy after an exam has done well on the basis of reasoning like this. Our puzzle is to explain what exactly is wrong with Tom's reasoning.

Case 2 is a variation on Dogramaci's (2010) and Weartherson's (2012) examples. But it is worth taking a moment to consider a salient difference between their examples and my own. My examples, unlike Dogramaci's and Weatheron's ones, do not employ an explicitly *statistical* form of inference.[3] Thus the fact that Tom's reasoning apparently exhibits just the same flaw as the reasoning in their examples suggests that the problem does not have to do with statistical reasoning in particular (Weatherson is clear that his conclusions are meant to apply to all ampliative reasoning). While I will freely use the term 'inductive reasoning' in connection to the puzzle, I will not presume that such reasoning is statistical: on the contrary, I think our puzzle cases are best handled by an alternative conception of the rational basis of ampliative reasoning. I return to this in Section 3.

Let us now turn to possible diagnoses. One quick response might be that nothing is wrong with Tom's reasoning. After all, one of his steps is inductive rather than deductive, and inductive reasoning is not guaranteed to be truth-preserving. It is to be expected that it will lead us astray sometimes. This response, however, is too quick. The problem with Tom's reasoning is not that it has a false conclusion — the problem is that it is, intuitively, *irrational*. And whether a piece of reasoning is rational does not simply reduce to whether the patterns employed in it are guaranteed to be truth-preserving: Bob's reasoning in Case 1 is not irrational, although it seems to employ the same potentially non-truth-preserving pattern as Tom's. What explains the difference?

Dogramaci's diagnosis is that inductive reasoning is incompatible with reasoning by universal generalisation. Reasoning by UG involves introducing an arbitrary object, reasoning about that

[3] Weatherson's (2012, 79) IR is a limiting case of statistical inference, as it only licences conclusions that have probability 1. I count it as statistical inference because it relies only on information about the distribution of a predicate *G* in a class of objects *F*.

object, and then generalising your conclusions over the entire class of objects that the original arbitrary object was meant to represent. Under what conditions is it rational to reason in this way? According to Dogramaci:

> The proper basis of rational reasoning by UG is recognition that the earlier reasoning is truth-preserving when the arbitrary object is replaced by any of the individual objects that it is standing in for. (Dogramaci 2010, 414-5)

Since step 2 in Tom's reasoning is not universally truth-preserving, Dogramaci's requirement would explain why Tom is not rational in generalising in step 4.

But Dogramaci's diagnosis does not go deep enough. Consider the following case, which, on the face of it, does not involve UG:

> *Case 3*. Fred begins by introducing the supposition that there exists someone who is happy after an exam despite having done poorly. Then he instantiates to an arbitrary individual *a*, who is happy after an exam despite having done poorly on it. By conjunction elimination, he infers that she is happy after her exam; and by inductive reasoning he infers that she did well on it. By another step of conjunction elimination from his supposition he gets a contradiction. By *reductio*, he concludes that there exists no person who is happy after an exam despite having done poorly on it.

This argument seems irrational in just the same way Tom's argument did, but appears to involve no use of UG. Our problem is not restricted to UG.

Dogramaci seeks to avoid this conclusion, by suggesting that a *reductio* argument proving the negation of an existential generalisation is 'logically' — though not 'nominally' — equivalent to reasoning by UG (Dogramaci 2010, 419–20). It is unclear what Dogramaci has in mind, however. For one thing, the *rules* used in the two arguments are clearly not equivalent: you cannot, for instance, prove the validity of *reductio* using just UG and conditional proof. Dogramaci might reply that our concern here is not with formal proof but with patterns of reasoning. Even so, it is hard to see in what sense the two patterns are equivalent. The only salient similarity between the two patterns seems to be that they both rely on reasoning about *arbitrary objects*. But those arbitrary objects are employed in very different ways: in Case 2 Tom *generalises* from the arbitrary case to a universal claim, while in Case 3 Fred derives a contradiction about *that particular* case. It is true that the case that Fred reasons about is arbitrary. But *reductio* does not require generalising from this case to a universal claim: indeed, given that the arbitrary case Fred reasons about is introduced by existential instantiation, such a step is not permissible anyway. Moreover, such a step is not needed: it is enough to show a hypothesis false if it leads to absurdity in just a *single* case. Of course, Fred's conclusion is equivalent to a universal generalisation; but this is because it is the negation of an existential claim, not because of a UG step in his reasoning.

Weatherson (2012) is also unsatisfied by Dogramaci's diagnosis: on his view inductive reasoning within suppositions should be altogether prohibited. But this seems like an overreaction. For one thing, Weatherson offers no explanation for *why* such reasoning is problematic. Moreover, there are examples of inductive reasoning within suppositions that seem rationally all right. Suppose that, just before seeing Raji with a big smile on her face, Bob reasons as follows:

> *Case 4*. Bob begins by supposing that Raji looks happy after her exam. On that supposition, he infers that Raji did well on her exam. On this basis, he forms the conditional belief that if Raji looks happy, she has done well on her exam.

Prima facie Bob's reasoning here seems rational, although it contains an inductive step within the scope of a supposition. At the very least it does not appear to suffer from the same blatant flaw present in cases 2 and 3. But Weatherson's account does not give us the resources to explain this difference.

So what, then, is the matter with the style of reasoning exemplified in cases 2 and 3? My suggestion is this: the problem is that subjects in these cases reason inductively about *arbitrary objects*.[4] In my view, what these cases highlight is that inductive reasoning about arbitrary objects is rationally prohibited.

But why exactly should reasoning about arbitrary objects be restricted in this way? My suggestion is based on two claims: first, that inductive reasoning is fundamentally reasoning about *what*

[4] Although Dogramaci does not discuss reasoning with arbitrary objects in detail, his view appears to be that if we take talk of arbitrary objects at face value (i.e., not as disguised quantifier talk), then we should find inductive reasoning about arbitrary objects unproblematic (Dogramaci 2010, 410–11). As I will argue, this is a mistake.

normally happens (in a non-statistical sense); and second, that in reasoning about the members of some class *F*, there is a crucial distinction between the *normal* and the *arbitrary F*. I discuss both of these claims in the next section. As we will see, my suggested diagnosis easily handles cases 2 and 3.

3. Reasoning and Normality

Return to Bob's reasoning in Case 1: intuitively, seeing Raji look happy immediately after taking an exam gives Bob good (though defeasible) reason to believe that she did well on her exam. But, of course, we can easily think of situations in which this connection fails (and so can Bob): Raji might, for example, have deliberately failed her exam in a fit of pique. So how can Bob conclude that Raji did well on her exam? My suggestion is that such situations are, given Bob's knowledge about people in general and Raji in particular, *abnormal*; and this, I suggest, is what entitles Bob to infer that Raji did well on her exam, even if such possibilities are not strictly inconsistent with anything that he knows.

This suggestion is in tension with the sorts of examples that Dogramaci (2010, 409) and Weatherson (2012, 79) use to introduce their puzzle, since they appeal only to *statistical* facts rather than facts about normality. I do not have a general argument to the effect that statistical reasoning of this sort — i.e., reasoning that concludes in *full belief* in non-probabilistic propositions based on merely statistical information — is never rational. But I think we have reason to believe that much ampliative reasoning proceeds on rather different grounds.

Consider the following cases.[5] When setting out for work in the morning you form a belief about your car's location based on your memory of where you parked it last night. This seems perfectly rational, even if you know that the incidence of car theft in your neighbourhood, although low, is not zero.[6] But our assessment of your reasoning seems to contrast with our assessment of familiar lottery cases. Suppose that you hold just one ticket in a

[5] The cases themselves are familiar from other debates in epistemology (Hawthorne 2004; Vogel 1990). Closer to present concerns is Smith's (2010, 19–22) discussion, although Smith is interested in the justification of beliefs rather than inferences.

[6] Of course, if car theft is *too* common in your neighbourhood we might hesitate to take the inference in question to be rational. But it is not obvious that it is the difference in statistical facts as such that makes the epistemic difference: perhaps what matters is that in such a situation your car's being stolen no longer seems abnormal.

large, fair lottery with just one winner. You cannot infer from your knowledge that your ticket *is very unlikely* to win to the conclusion that your ticket *will not* win — with the result that, e.g., you throw it away as worthless. Moreover, this asymmetry remains even if we stipulate that your chance of winning the lottery is no higher than the probability that your car was stolen overnight.

The hypothesis that inductive reasoning is reasoning about what normally happens — rather than about what *usually* happens — neatly explains this asymmetry.[7] While your evidence is consistent with your car's not being where you parked it last night, this possibility is *abnormal* (cars do not vanish into thin air, nor do they drive off by themselves), and so does not undermine your inference. By contrast, your lottery ticket's turning out to be a winner would not be abnormal. Since you know that the lottery is fair, you know that your ticket's winning is just as normal an outcome as any other ticket's winning. So what reason could you have for setting aside the possibility that *your* ticket will be the one that wins?

Now, as Vogel (1990, 21-22) notes, we can construe the car theft case in such a way that our judgments about it parallel those in the lottery case. But it is instructive to see just *how* Vogel achieves this:

> There is reason to think that some car or cars similar to your will be stolen, and you have no non-arbitrary ground for believing that your car won't be the one (or among the ones) stolen. (Vogel 1990, 22)

In effect, Vogel highlights a sense in which your car's being stolen is *not* abnormal: assuming that *some* cars relevantly like yours, chosen at random, were stolen overnight, there is nothing abnormal about *your* car's being among them. But then the present proposal has no trouble accounting for the shift in our intuitions: if there is nothing abnormal about your car's having been stolen overnight, you cannot set aside the possibility that it was.

[7] An alternative reaction might be that in the lottery case you fail to *know* that your ticket will lose — and it is this, rather than any irrationality in your reasoning, that explains why throwing away your ticket seems irrational. But this seems strange. If your reasoning is rationally fine, and your ticket really is (let us assume) a loser, then why should your conclusion fail to be knowledge? On these assumptions the lottery case would be a Gettier case; but it seems to lack any of the usual Gettier funny business.

As Vogel and others (Hawthorne 2004; Smith 2010) have noted, the fact that we can think about such cases in two strikingly different ways raises a puzzle: taking all your evidence into account, is your judgment as to the whereabouts of your car this morning justified or not? It is important to note, however, that this puzzle is not ours. From the present point of view, what the shift in our intuitions shows is just that while we are originally prone to assuming that it would be abnormal for your car to vanish overnight, there is a relatively easy way of getting us to abandon this assumption; and there is nothing puzzling about the idea that, *given* different background assumptions, different conclusions might be rationally available to you. To decide whether your judgment that your car is where you parked it last night is, all things considered, epistemically justified would require (among other things) deciding *whether the assumption in question is itself warranted* on your evidence; and while I will have more to say about normality in what follows, I do not expect what I say here to settle this question.[8]

Setting purely statistical inference aside (to the extent that there is such a thing), I think that this approach to ampliative reasoning has very broad application. It is a familiar idea that ampliative reasoning works by exploiting patterns and regularities in our environment, even though those patterns and regularities are defeasible, or admit of exceptions. Such regularities might be expressed by conditionals, or *ceteris-paribus* or 'default' rules. The question, then, is how it can be rational to base one's reasoning on such regularities. And the answer, I suggest, is that we treat the exceptions as *abnormal*, and *idealise away* from them (see also Smith 2010, 16).

It is perhaps more traditional to think of inference to the best explanation as the alternative to statistical inference (Harman 1973). However, as I will explain below, normality in the sense intended here is also to be understood in explanatory terms. Thus the two approaches need not compete with each other: the appeal to normality can be seen as a way of codifying the intuitions behind inference to the best explanation (Boutilier 1994, for example, sees it this way).

[8] Notice, incidentally, that these considerations might allow us to explain away whatever intuitive appeal purely statistical rules have ('every normal person', Dogramaci 2010, 409, thinks, accepts his statistical rule). While high statistical frequency is not the same thing as normality, we might have a tendency to take it to be *a sign of* normality: if we know that 99.999% of the *Fs* are *G*, we are prone to assume that it is *normal* for the *Fs* to be *G*.

We can put all this a little more formally, as follows. Begin with the familiar idea that a subject's beliefs determine a set of worlds that are *epistemically open* to her: these are the ways things might be, given what she believes. We add, moreover, that our subject's existing beliefs impose an *ordering of normality* over worlds: given any two worlds, our subject can rank them in terms of normality.[9] Then the suggestion is this: if our subject reasons from premiss *P*, she can infer *Q* just in case *Q* is true in *all the most normal P-worlds*, given what she antecedently believes.

How should we understand the relevant notion of normality? While there is, of course, a statistical notion of normality, that is not what is at issue here. Normality in our sense is a matter of things happening the way they *should*. But it is important not to read too much into this idea. Suppose Mary reasons in accordance with the defeasible rule that if something is a mammal, then it does not lay eggs. From the present point of view, this implies that Mary treats platypuses and echidnas as in some sense abnormal, *qua* mammals. But what this means is just that Mary treats them as *exceptions to certain regularities that she relies upon* in navigating the world, and which are, as such, in need of some special explanation (for example, some story about how early the monotremes branched off from the main line of mammalian evolution). It does not imply that she must treat them as abnormal in some absolute sense.

Following Smith (2010), I suggest we can spell normality out in terms of the *explanatory regularities* that we take to hold in our environment. Such regularities are generally not exceptionless laws, but rather hold only *ceteris paribus*, or other things being equal. To say that an object or event of kind *F* is abnormal is to suggest that some of its features are *exceptional*, and thus require special explanation, over and above the regularities that we take to govern the *Fs* in general. We can also compare states of affairs and even entire worlds as to normality: as suggested earlier, we seem to think that a world in which your car in the morning is where you parked it last night is more normal than a world in which it is not. Such assessments will be complex and potentially context-dependent. A world might

[9] Here I follow Boutilier (1994). As Boutilier argues, any reasonable normality ordering is *transitive* and *reflexive*. Worlds consistent with what the subject believes are maximal elements in this ordering. This does not mean that the subject believes the world to be maximally normal in some absolute sense; it simply reflects the fact that we are modelling normality *given what the subject already believes to be the case*. I also assume that for each *P* that is true in any world, there is a most normal set of worlds in which it is true.

conform very well to certain regularities but only at the expense of grossly violating others. Which of two (or more) competing sets of regularities we privilege in our assessments of normality depends on which we take to be more explanatorily fundamental.

Now, these brief remarks obviously fall short of a full account of normality in the relevant sense. Questions arise: will it always be possible to rank worlds in terms of normality, independently of our intuitions about the rationality of associated inferences? And, to what extent (if any) will such rankings depend on contextual factors? These questions must await a different occasion. For present purposes, what matters is just the structural point that in inductive reasoning we exploit defeasible regularities by *idealising away* from exceptions. This, as we shall see, suffices for the task immediately at hand — namely, explaining what goes wrong in the puzzling cases we encountered in Section 2.

Note, first, that my suggestion clearly allows for inductive reasoning within suppositions. Even before seeing Raji with a big smile on her face, Bob can focus his attention to possible worlds in which Raji looks happy after her exam; and he can conclude that, in the most normal among them, she has done well on it. The problem highlighted by Dogramaci's examples arises only in the case of reasoning inductively about *arbitrary objects*. Bob can give a determinate answer to the question whether the most normal world in which *Raji* looks happy after an exam is a world in which she has done well on it. But Tom can give no determinate answer to the corresponding question about the *arbitrary person*.

This is evident on Kite Fine's (1985) account of the semantics of reasoning with arbitrary objects. On Fine's view, the 'instantial terms' used in reasoning refer to special 'A-objects'. A-objects are distinct from all ordinary objects, but are associated with value-ranges of such objects. So how do we reason about them? According to Fine, 'the basic principle is that a sentence concerning A-objects is true (false) just in case it is true (false) for all of their values' (Fine 1985, 41). Thus, A-objects are only determinate in respects in which *all objects in their value range agree*. But then, since the generalisation that people who look happy after an exam have done well on it is not an exceptionless law, there simply is *no* determinate answer as to whether the most normal worlds in which the *arbitrary* person looks happy after an exam is a world in which she has done well on it.

As Breckenridge and Magidor (2012) argue, there may be reasons to be dissatisfied with Fine's account and especially with his commitment to special A-objects. On their account, instantial terms refer to ordinary, fully determinate individuals — although it is both arbitrary and unknowable *which* individual each such term refers to. This is no problem for my account, as the ban on inductive reasoning with arbitrary objects does not depend on the existence of special A-objects; it depends, rather, on the role appeal to arbitrary objects plays in reasoning. There is a sense in which appeals to the normal case and to the arbitrary case both function in reasoning as *representatives* of some class of entities. But they do so in very different ways. A noted earlier, the point of reasoning about the normal case is that it allows us to exploit regularities and patterns in our environment by focusing on what things should be like — i.e., by idealising away from exceptions. By contrast, when we reason about the arbitrary F no idealisation is involved: instead of ignoring exceptional Fs we *ignore any properties not shared* by all Fs. And this is precisely the norm that both Tom and Fred violate in cases 2 and 3: even though the regularity that people who look happy after an exam have done well on it is not an exceptionless law, they apply it to the *arbitrary* person who looks happy after an exam. This, rather than the application of UG, is the real source of the problem.

Now, it is of course a consequence of this view that UG is rationally incompatible with inductive reasoning. But it is worth noting that this does not entail that we cannot reason inductively for general conclusions. Suppose that you have observed a large number of emeralds (identified by their chemical and microstructural properties) and noticed that they all share a distinctive green color. Assuming that you have a reasonable set of background beliefs, it seems plausible that worlds in which *all* emeralds are green will be more normal for you than worlds in which some are green and some are not. These latter worlds involve physical objects that are alike in their chemical and microstructural properties but different in their macroscopic properties. What accounts for such differences? It seems clear that some sort of special explanation would be required. The present account, therefore, predicts that it would be rational for you to infer that all emeralds are green.

It is clear that whether such general conclusions are available to a subject will depend on her background beliefs. While worlds in which objects that share their chemical and microstructural properties also share their macroscopic ones are plausibly more normal than worlds in which this is not the case, worlds in which everyone who looks happy after an exam has done well on it are *not* plausibly more normal than worlds in which this is not so. Given what we know about the variety of human motives, such worlds look bizarre: do people there only care about exam results? Thus we cannot conclude that everyone who looks happy after an exam has done well on it.

I have argued that the puzzle cases Dogramaci (2010) brought to our attention are best addressed by recognizing that we cannot reason inductively about arbitrary objects. This restriction, in turn, flows naturally from thinking about inductive reasoning as reasoning about what normally happens. As I have tried to show, moreover, this is a view of inductive reasoning that has considerable independent plausibility.

University of New South Wales
m.valaris@unsw.edu.au

References

Bird, Alexander. 2005. Abductive Knowledge and Holmesian Inference. *In Oxford Studies in Epistemology 1*, edited by Tamar Gendler and John Hawthorne, 1–31. 1. Oxford University Press.

Boutilier, Craig. 1994. Unifying Default Reasoning and Belief Revision in a Modal Framework. *Artificial Intelligence* 68 (1): 33–85.

Breckenridge, Wylie, and Ofra Magidor. 2012. Arbitrary Reference. *Philosophical Studies* 158 (3): 377–400.

Dogramaci, Sinan. 2010. Knowledge of Validity. *Noûs* 44 (3): 403–432.

Fine, Kit. 1985. *Reasoning with Arbitrary Objects*. Oxford: B. Blackwell.

Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press.

Hawthorne, John. 2004. *Knowledge and Lotteries*. New York: Oxford University Press.

Smith, Martin. 2010. What Else Justification Could Be. *Noûs* 44 (1): 10–31.

Vogel, Jonathan. 1990. Are There Counterexamples to the Closure Principle? In *Doubting: Contemporary Perspectives on Skepiticism*, edited by Michael David Roth and Glenn Ross, 13–29. Dordrecht: Kluwer.

Weatherson, Brian. 2012. Induction and Supposition. *The Reasoner* 6: 78–80.