ARE REASONS EVIDENCE OF OUGHTS?

Franck LIHOREAU

ABSTRACT: In a series of recent papers Stephen Kearns and Daniel Star argue that normative reasons to ϕ simply are evidence that one ought to ϕ, and suggest that "evidence" in this context is best understood in standard Bayesian terms. I contest this suggestion.

KEYWORDS: reasons, evidence, oughts, normativity, rationality

Reasons as Positively Relevant

What is it for a reason to have a certain strength? And what is it to weigh a reason against another? According to Stephen Kearns and Daniel Star,[1] if we accept the following claim about normative reasons for action:

> (RA) Necessarily, a fact *F* is a reason for an agent *A* to ϕ iff *F* is evidence that *A* ought to ϕ (where ϕ is an action)[2]

an answer to those questions can be fleshed out in terms of a familiar concept of evidence on which we already have a good, independent grasp, namely the concept of "incremental" evidence understood in standard Bayesian terms (or "positive relevance"):

> (IE) *E* is evidence for *H* iff $Pr(H \mid E) > Pr(H)$

that is, iff the probability of *H* when *E* is added (to one's prior background information, as encapsulated by *Pr*) is strictly greater than the probability of *H* (on one's prior information) alone. By combining this definition of evidence with the general claim about reasons in (RA), we get straightforward answers to the questions we started with: the strength of a reason to ϕ is the degree to which it

[1] Daniel Kearns and Kenneth Star, "Reasons: Explanations or Evidence?" *Ethics* 119 (2008): 31-56; "Reasons as Evidence," in *Oxford Studies in Metaethics 4*, ed. R. Shafer-Landau (Oxford: Oxford University Press, 2009): 215-242; "Weighing Reasons," forthcoming in *Journal of Moral Philosophy*; "Reasons, Facts-about-Evidence, and Indirect Evidence," forthcoming in *Analytic Philosophy.*

[2] Kearns and Star, "Reasons as Evidence," 216.

raises the probability that one ought to ϕ, and the stronger the reason to ϕ, the more probable it is that one ought to ϕ.[3]

My purpose in this note is not so much to contest the otherwise intriguing "reasons as evidence" thesis defended by Kearns and Star as it is to cast doubt on the suggested appropriateness of a standard Bayesian understanding of evidence for making sense of the "strength" and the "weighing" of reasons for action. To this end, I offer two counterexamples to the idea that all reasons to ϕ increase the epistemic probability that one ought to ϕ,[4] thereby establishing that within the scope of the account of incremental evidence in (IE), the account of normative reasons for action in (RA) is too narrow: not all reasons to ϕ are evidence that one ought to ϕ. As we shall see in due course, this result carries over to other standard Bayesian accounts of evidence as well.

Evidentially Irrelevant Reasons

A reason can fail to raise the (epistemic) probability, and therefore, assuming (IE), to be evidence that one ought to do something – or so shall I argue. This situation can arise in two different ways.

First, when the reason simply is "evidentially irrelevant" to the corresponding ought: A fact *F* can be a reason to ϕ even when *F* does not affect – neither raises nor lowers – the probability that one ought to ϕ.

To see this, consider the following case (directly inspired by, albeit freely adapted from an example by Peter Achinstein[5]).

Suppose you enjoy drinking a certain soda so much that you usually buy it by batches of 100 bottles; and today, you drank one and only one bottle of that soda from such a batch – call it batch *b* –, and no one else did. Now, consider the following claims about *b*:

> (*E1*) Newspaper 1 reports that 99 out of the 100 bottles in b are contaminated by an extremely dangerous and highly contagious virus.
>
> (*E2*) Newspaper 2 makes the same announcement as Newspaper 1 about *b*.
>
> (*H*) You have drunk from a contaminated bottle.

[3] Kearns and Star, "Reasons as Evidence," 232. See also Kearns and Star, "Weighing Reasons."

[4] In contesting this idea, I side with John Brunero ("Reasons and Evidence One Ought," *Ethics* 119 (2009): 538-545). But the lesson I draw differs from his in scope and is somewhat less categorical. For a discussion of Brunero's arguments, see Kearns and Star, "Weighing Reasons."

[5] Peter Achinstein, "A Challenge to Positive Relevance Theorists: Reply to Roush," *Philosophy of Science* 71 (2004): 521-524.

Clearly, E2 is not evidence for H since it neither increases nor decreases the probability of H:

$$Pr(H \mid E2 \,\&\, E1) = Pr(H \mid E1) = .99.$$

Because (the contents of) the reports are the same, adding Newspaper 2's report does not make it more probable that you've drunk from a contaminated bottle than does Newspaper 1's report alone.

Now, as a matter of public health, your drinking from a bottle contaminated by an extremely dangerous and contagious virus creates an obligation for you to put yourself into quarantine and stay at home; and although there might arguably be exceptions, the probability that you ought to do so remains nonetheless a strictly increasing function of the probability of *H*. Therefore, since *E2* does not affect the probability of *H*, it does not affect the probability that you ought to put yourself into quarantine either. So, if (IE) is true, *E2* is not evidence that you ought to put yourself into quarantine (and is even evidentially irrelevant to such an obligation).

However, it should be clear and uncontroversial that the .99 probability of *H* is more than enough, given *E1* (and the relevant public health obligations), for *E2* to be a reason – and a good one – for you to keep yourself in quarantine. So, the probability that you ought to keep yourself in quarantine is not affected by *E2*, despite the fact that, given *E1*, *E2* is a reason for you to keep yourself in quarantine.

Now, it is not strained to think that it is not properly speaking the newspaper report, but rather your drinking from a possibly contaminated bottle, that is a reason here. But Kearns and Star cannot afford this thought. For they explicitly state that a report or announcement of a fact – e.g. a newspaper report to the effect that people are starving in Africa – "has just as good a case for being a reason [viz. to send money to Oxfam] as do more paradigmatic reasons (such as that people are starving in Africa)."[6]

Moreover, note that mentioning the possible unreliability of the newspapers would be irrelevant here. Whether we assume the reports to be truthful or not, the result is the same. Let *N1* be that Newspaper 1 tells the truth about batch *b* and virus *v*, and *N2* that Newspaper 2 also tells the truth about *b* and *v*. Since as a matter of fact $Pr(H \mid E2 \,\&\, N2 \,\&\, E1 \,\&\, N1) = Pr(H \mid N2 \,\&\, E1 \,\&\, N1)$, *E2* fails again to raise the probability of *H*, and therefore of your obligation to put yourself into quarantine, despite the fact that *E2* still remains a reason for you to put yourself into quarantine (given *N2* & *E1* & *N1* this time). Therefore, the question of the newspapers' reliability does not arise here.

[6] Kearns and Star, "Reasons, Facts-about-Evidence, and Indirect Evidence," msp 1; see also "Reasons as Evidence," 233-234.

So, it seems that a fact can be a reason for one to do an act even if, because it does not affect the probability that one ought to do this act, the fact is not evidence that one ought to do the act. A consequence would be that assuming (IE) as an account of evidence, the analysis of reasons in (RA) is too narrow.

An easy fix would be to suggest replacing > with the weaker ≥ in the "naïve" analysis of incremental evidence we started with:

(IE*) *E* is evidence for *H* iff $Pr(H \mid E) \geq Pr(H)$.

This suggestion – at odds with standard Bayesian approaches to evidence – would fix the problem in the case at hand, since the condition on the right-hand side would *ipso facto* be satisfied. But a more serious problem is lurking.

Negatively Relevant Reasons

Failure of a reason to raise the (epistemic) probability that one ought to do something can indeed stem from another source as well, namely, from the reason being "negatively relevant" to the corresponding ought. In other words, a fact *F* can be a reason to ϕ even when, instead of raising it, *F* lowers the probability that one ought to ϕ.

To see this, take the following example. You own a restaurant that serves exotic food, and some highly perishable good, *g*, is being shipped to you as part of a bigger batch of perishable goods. Consider the following claims about *g*:

(H) *g* has gone off.

(E1) the shipping takes *n* days.

(E2) *g* is shipped as part of batch *b*.

And suppose you have somehow determined that good *g* has a 90% chance of having gone off if the shipping takes *n* days:

(1) $Pr(H \mid E1) = .9$,

that 75% of the goods in batch b have already gone off by the time they are sent:

(2) $Pr(H \mid E2) = .75$,

that the shipping of good *g* as part of batch *b* has a very low .088 probability of taking n days:

(3) $Pr(E2 \,\&\, E1) = .088$,

and that there is a not much higher .075 probability that good *g* has gone off after a *n*-day shipping inside batch *b*:

(4) $Pr(H \,\&\, E2 \,\&\, E1) = .075$.

Then, by the definition of conditional probability, we get an 85% chance that good *g* has gone off if it was shipped in *n* days inside batch *b*:

(5) $Pr(H \mid E2 \,\&\, E1) = Pr(H \,\&\, E2 \,\&\, E1) \,/\, Pr(E2 \,\&\, E1) \cong .85.$

As a consequence, the probability of *H* is actually lowered by *E2*, since:

(6) $Pr(H \mid E2 \,\&\, E1) \cong .85 < Pr(H \mid E1) = .9.$

Now, as a matter of public health and food regulations, the circumstances in which a restaurant might be allowed to serve spoilt food to its customers are presumably very, very few. I believe Kearns and Star will see no objection in conceding that conditional normative principles exist whereby if something like the nonnormative fact that the food is spoilt obtains, then so does something like the unconditional normative fact that one ought to throw it away and not serve it.[7]

And the probability that you ought to throw g away and not serve it to your customers is thus presumably a strictly increasing function of the probability of *H*. So, the probability that you ought to throw *g* away too is presumably lowered by *E2*. So, if (IE) is true, *E2* is not evidence that you ought to throw *g* away (and is even evidence that you ought not to throw *g* away).

But *E2* is a reason to throw *g* away. This point is relatively unproblematic: not only does the fact that *g* was part of *b* constitute a reason to throw it away, it constitutes a *very good reason* to do so by most people's standards given that 75% of the goods coming from batch *b* have already gone off by the time they are sent. So, the probability that you ought to throw *g* away is lowered by *E2*, despite the fact that, given *E1*, *E2* is a reason for you to throw *g* away.

Therefore, a fact can be a reason for one to do an act even if, because it lowers the probability that one ought to do this act, the fact is not evidence that one ought to do the act. As a consequence, the left-to-right reading of (RA) is false within the scope of (IE) – no need to say that it is false within the scope of (IE*) as well.

Discussion

It goes without saying that our proposed counterexamples to (RA) are not isolated examples, that it is not difficult to generate myriads of structurally similar examples, and that many other probability assignments could have been used to reach the same conclusion that some reasons are not evidence for oughts. Also, if such examples constitute genuine cases of practical reasons, these cases are "standard" in that "they are examples where the relevant facts are transparent to the agent, that is, where

[7] cf. Kearns and Star, "Reasons as Evidence," 229.

there are no false beliefs playing any role in deliberation and there is no misleading evidence around clouding the water."[8]

Still, one could respond to our counterexamples in one of two ways.

The first is simply to ignore the intuitive pull we undeniably feel towards considering them genuine cases of practical reasons at all and somehow insist that there are not. However, this line of response is not available to Kearns and Star since those cases satisfy the various sufficient conditions they state for being cases of practical reasons. They explicitly defend that if a fact *F* "can play an appropriate role in one's reliably concluding that one ought to ϕ",[9] or if it "can play an appropriate public role in rationally convincing [someone] that she ought to ϕ and in rationally convincing other people that she ought to ϕ",[10] or else if "it is normally the case that if a fact relevanty similar to *F* obtains, then one ought to do something relevantly similar to ϕ-ing",[11] then that fact *F* is a reason to ϕ. But it is uncontroversial that a newspaper report on a contaminated batch of goods as in the first of our cases, or a product being part of such-and-such batch of merchandise as in the second case, typically constitute information that can help us determine what we ought to do, help us justify what we do, convince others about what they ought to do, and enter into the formulation of normative principles connecting relevantly similar information with obligations towards relevantly similar actions (as reflected in health conventions and food regulations, for instance). So, Kearns and Star will have to concede that the relevant facts involved in our putative counterexamples are indeed reasons to do the relevant acts.

The second way one could respond to our cases is to opt for a different Bayesian account of evidence. Among some of the other relatively standard options available, one is to drop the incremental notion of evidence in (IE) in favor of an "absolute" one:

(AE) *E* is evidence for *H* iff $Pr(H \mid E) > k$, for some degree *k* of high probability

while another option is to go for a "probative" notion of evidence instead:

(PE) *E* is evidence for *H* iff $Pr(H \mid E) > Pr(H \mid \neg E)$.

Unfortunately, none of these options will work.

In (AE) the appropriate threshold k can undoubtedly be set very high. Still, in the context of the *reasons as evidence* thesis, it will have to be set low enough to

[8] Kearns and Star, "Reasons as Evidence," 223.
[9] Kearns and Star, "Reasons as Evidence," 225.
[10] Kearns and Star, "Reasons as Evidence," 227.
[11] Kearns and Star, "Reasons as Evidence," 228.

do justice to the fact that, by most people's standards, the facts involved in our purported counterexamples do count as reasons, and even good reasons to do the relevant acts: an 85% conditional probability that the food has gone off and a 99% conditional chance of having drunk from a contaminated bottle are, in this respect, more than high enough. So, the proposed examples are counterexamples to (RA) within the scope of (AE) too.

As to (PE), in the first of our cases the epistemic probability of having drunk from a contaminated bottle when Newspaper 2's report is added is not higher than the epistemic probability of having done so in the absence of Newspaper 2's report, since this probability already is 99% given the report made by the other newspaper, Newspaper 1. In our second case the epistemic probability that the food is not part of the incriminated batch b and takes n days to be shipped (i.e. $\neg E2$ & $E1$) can easily be specified so that the epistemic probability that the food has gone off if it was part of this batch and took that long to be shipped is lower than the epistemic probability that it has gone off if it was not part of that batch yet took that long to be shipped. So, the objection that stems from our counterexamples carries over from (IE) to (AE) and (PE).

Conclusion

To sum up, the probabilistic implementation that Kearns and Star suggest to put flesh on the bones of their general *reasons as evidence* thesis is too narrow: (RA) fails to provide a necessary condition for being a normative reason for action within the scope of the Bayesian account of incremental evidence they suggest, (IE), and this is true as well with other standard Bayesian understandings of evidence, like the account of absolute evidence in (AE) and the analysis of probative evidence in (PE).

It would certainly not be fair to conclude from this to the inadequacy of (RA) itself. As Kearns and Star remark, the claim that results from combining (RA) with (IE) is more specific than the claim in (RA), and likewise with (AE) and (PE). But they do place hopes in the possibility of explaining what is involved in weighing reasons against each other in terms of a particular account of weighing evidence along standard Bayesian lines.

So what we may conclude from the foregoing considerations is that appeal to such standard Bayesian accounts of evidence as those in (IE), (AE), or (PE) seems inappropriate to Kearns and Star's general purpose of making sense of the notion of

a reason's strength and of weighing reasons. Their hopes in this respect might not be so well-placed as they think.[12]

[12] An earlier version of this material was presented at the New University of Lisbon. I wish to thank the audience for many helpful comments. My work on this paper was carried out in part within the "Argumentation, Context, and Communication" project (PTDC/FIL-FIL 110117/2009) supported by the Portuguese Foundation for Science and Technology.