

ISSA Proceedings 2010 - On The Priority Of Epistemic Rationality



1. Introduction

One influential way to think about arguments is the following: an argument consists of premises asserted in support of an asserted conclusion; the purpose of arguments is to rationally persuade their audience of the truth of their conclusions; good arguments are those that achieve their purpose. On this picture, in order for an argument to achieve rational persuasion, its premises must be rationally acceptable to the participants in the argument, and it must be rational to think that the premises support the conclusion. And, if we take the type of rationality relevant to the assessment of arguments to be epistemic rationality, then the theory of epistemic rationality becomes directly relevant to the theory of argument.

What I want to do in this paper is to try to show that epistemic rationality is not a matter of believing in such a way as to maximize the likelihood of achieving our epistemic goals. If it were, then epistemic rationality would be a species of practical rationality. But it cannot be a species of practical rationality, because it is prior to practical rationality. It follows that epistemic rationality is not a matter of achieving our epistemic goals.

In the context of the theory of argument, it is particularly important to see that epistemic rationality is not a matter of believing so as to achieve our epistemic goals: if it were, then for an agent who does not care about achieving an epistemic goal, nothing would count as epistemically rational or irrational. It would then follow that for a subject who lacked an epistemic goal, no arguments could count as good or bad. An epistemic approach would have nothing to say about arguments in such cases. I take it that that would be the wrong result, and a serious mark against the epistemic approach to argument evaluation, because the goodness or badness of arguments should not depend on whether people have an epistemic goal.

2. Epistemic rationality

The dominant way to think of epistemic rationality is in teleological terms. The standard picture is that we have an epistemic goal, and epistemically rational

beliefs are those that achieve (or those that we would, on reflection, take to achieve) our epistemic goal; beliefs that fail to achieve our epistemic goal are epistemically irrational. There are various ways to specify the content of our epistemic goal; most epistemologists pick up on William James' idea that epistemic rationality is about achieving true beliefs and avoiding errors. William Alston (1985), for example, holds that the epistemic goal is to maximize truth and minimize falsity in a good-sized body of beliefs. Other views are that our epistemic goal is to believe all of the truths that there are, and nothing else (Latus 2000), to have true beliefs and not to have any false beliefs right now (Foley 1987), to maximize truth and minimize falsity in a large body of beliefs over an extended period of time (Vahid 2003), or that we have a variety of epistemic goals, such as truth, justification, knowledge, simplicity, etc. (Kvanvig 2005).

Now, most epistemologists do not specify whether subjects must in fact have the epistemic goal, in the sense that it must be something that they want to achieve, in order for their beliefs to count as epistemically rational when they do achieve it, or as epistemically irrational when they fail to achieve it. Following Kelly (2003), we can call the position that makes epistemic rationality a matter of believing in such a way as to achieve the epistemic goal that agents have (i.e. care about achieving) the "instrumentalist conception" of epistemic rationality (ICER). Foley (1987) is the clearest exponent of a developed account of epistemic rationality who accepts ICER. Robert Nozick (1993) also appears to accept ICER. Some theorists (e.g. David 2001) explicitly want to avoid ICER, while still maintaining that epistemic justification is a matter of achieving the epistemic goal. But many theorists simply do not commit one way or the other.

The alternatives to ICER are either to hold that epistemic rationality is a matter of believing in such a way as to achieve the epistemic goal, whether want to achieve it or not (a broader sort of teleological conception of epistemic rationality), or else to hold that epistemic rationality just has nothing to do with achieving a goal. Adopting the Kantian terminology, we can call these alternatives categorical conceptions of epistemic rationality, because they hold that epistemic rationality is independent of what people desire.

I don't think that ICER is correct. For my purpose here, it is not important which of the other two alternatives to take up, although it seems to me that the two main arguments in this paper undermine any kind of goal-directed conception of epistemic rationality. Whether they do so is not important for now, though; all

that is important to see is that the epistemic rationality of our beliefs does not depend on the content of the epistemic goals that we want to achieve. If it did, then for an agent who lacks an epistemic goal, nothing would count as epistemically rational or irrational.

Now, before we move on, I should make it clear that in what follows, I have in mind a very narrow conception of practical rationality, essentially treating practical rationality and instrumental rationality as the same. Everyone at least agrees that instrumental rationality is one important type of practical rationality. Some theorists stop there (e.g. Bertrand Russell, Larry Laudan, Herbert Simon, Richard Foley), and take instrumental rationality to be all that there is to rationality. Others take practical reason to encompass also the determination of what goals we ought to adopt, what forms of practical maxims are permissible, etc. I am stipulating here that practical rationality is instrumental rationality. Nothing important hangs on that stipulation here. If practical rationality is more than instrumental rationality, then the arguments in this paper can be recast to accommodate that. I make the stipulation just in order to keep things simple.

3. Two arguments against ICER

There are two arguments that I would like to bring to bear here against ICER. If epistemic rationality is a matter of believing so as to achieve our epistemic goals, then epistemic rationality is a type of practical rationality, because achieving our goals falls under the purview of practical rationality. The first argument here is intended to show that epistemic rationality is not instrumental in nature, because epistemic goals can be achieved in epistemically irrational ways. The second argument is intended to show that practical rationality depends on epistemic rationality.

3.1. Achieving our epistemic goal in epistemically irrational ways.

The first argument against ICER to consider is the fact that we can hold epistemically irrational beliefs that nevertheless promote the achievement of our epistemic goal. Consider a typical formulation of the goal: to have a favourable truth-falsity ratio in a good-sized body of beliefs.

Given a diachronic understanding of that goal, it is easy to construct examples of epistemically irrational beliefs that serve to achieve it. Consider, for example, a student who, contrary to all the evidence provided to her by her poor academic record in high school, believes in her academic ability, which gives her the

confidence required to study hard, score well on her SATs, get into a good college, and acquire all sorts of interesting true beliefs. She holds her belief in her academic ability against the available evidence, so it hardly counts as an epistemically rational belief, and yet it helps her to achieve her epistemic goal: even though it is a false belief, and she holds it against the evidence, it helps her to get into a good college and acquire all sorts of interesting true beliefs. Feldman puts the point nicely: “if believing something now would somehow lead me to believe lots of truths later, that long-term epistemic benefit is ... irrelevant to [the judgment of whether p is true]” (1988, pp.249-50).

Precisely in order to avoid this sort of problem, Foley (1987) makes his formulation of the epistemic goal synchronic: the epistemic goal, for Foley, is to believe all and only truths, right now. The purpose of this restriction is to screen off epistemically irrelevant factors from our doxastic deliberations and epistemic evaluations. A subject’s belief is epistemically rational on Foley’s account iff, after sufficient reflection, the subject would take the belief to satisfy the epistemic goal. No quantitative amount of reflection can be specified for reflection to be sufficient; sufficient reflection is just reflection to the point of reflective stability, so that further reflection would not lead the subject to change his mind.

Given this way of setting up the epistemic goal and what is required to achieve it, Foley takes it that only “uncontroversial” beliefs can satisfy it. A belief is uncontroversial for a subject, roughly, when the subject has available to him an argument that, upon sufficient reflection, he would take to support the truth of the belief **[i]**. Now, Foley’s notion of uncontroversiality provides us with a plausible account of what it takes to be epistemically rational, I take it, which is why it is important for him to be able to show that the set of beliefs that are uncontroversial for a subject, and the set of beliefs that the subject would take on sufficient reflection to satisfy the epistemic goal, turn out to be one and the same. In order to press the objection to instrumentalism, then, what we need is a counterexample to show that these two sets of beliefs do not turn out to be the same.

So what we need is a case of a belief that satisfies the epistemic goal, but fails to be uncontroversial for a subject, or else a belief that is uncontroversial for a subject but fails to achieve the epistemic goal. Both kinds of case can be constructed, I imagine, but I’ll only give an example of the first. The point of this case is to show that even the synchronic epistemic goal can be satisfied in an

epistemically irrational manner; restricting the epistemic goal this way fails to screen off epistemically irrelevant factors. (This should not be surprising, by the way; deviant ways of achieving ends are nothing new in philosophy.)

With all of that in mind, let us turn to the problem case. Suppose there is an agent, Larry, who takes himself to be infallible with respect to a domain of knowledge *D*, and he has taken himself to be so for some time. He has been mistaken on a few occasions, but he has successfully put those occasions from his mind. It pleases him to think that he is infallible, and he manages not to think about his few failures (like many of us, he is quite capable of ignoring evidence), so he continues to believe in his infallibility. Since the time when he formed the belief in his infallibility, Larry has produced very many beliefs within *D*, and he continues to hold those beliefs. Furthermore, he is aware of several scientific studies which agree that people who take themselves to be infallible with respect to domain *D*, for whatever reason, produce very many true beliefs about it, and very few false beliefs. The ratio of true to false beliefs, moreover, is much higher for people who believe themselves to be infallible than for those who do not. Finally, the studies also show that people who for whatever reason give up the belief in their infallibility also give up all their beliefs about *D*.

These studies do not, of course, figure in Larry's reasoning when he produces beliefs about *D*, because he believes himself to be infallible, so he does not need the extra boost to his epistemic self-confidence. But the studies do support his belief in his infallibility, in the following way. Larry has read and been impressed by Foley's book, and he wants to make sure that he is epistemically rational in his beliefs. He therefore proceeds to test his beliefs for how well they promote the epistemic goal of now having true beliefs and now not having false beliefs. He recognizes (because he has read the scientific studies to this effect) that because he takes himself to be infallible, he must have produced very many true beliefs and very few (if any) false ones about *D*. He concludes that his belief in his infallibility is an effective means to achieving the epistemic goal. He does not even bother to determine whether he has an uncontroversial argument in favour of his infallibility, because the belief just obviously promotes the epistemic goal. Even though it is in fact both false and controversial for him - since he has been mistaken on occasion, and he could make himself aware of his mistakes, if he reflected carefully - it promotes the epistemic goal so well, because it is only one false belief that allows him to hold many true beliefs. (Larry does not, of course,

take his belief in his infallibility to be false, but he can see that even if it was false, it would still clearly promote the epistemic goal, so he does not go on to wonder about its uncontroversiality.)

Because Larry succeeds in achieving the epistemic goal, his belief in his infallibility counts as epistemically rational, if ICER is true. But his belief in his infallibility ought to be obviously epistemically irrational, because it is held contrary to some conclusive available evidence that is being ignored. It even fails Foley's own test for epistemic rationality: it is not an uncontroversial belief, because upon a little serious reflection, Larry would see that he has good reason to doubt his infallibility. So, even though this formulation of the epistemic goal is designed to screen off epistemically irrelevant ways of achieving it, we have here a case of an epistemically irrational belief that nevertheless achieves it **[ii]**. And, although that is not a conclusive reason for rejecting ICER, it ought to undermine much of its appeal: ordinary conceptions of the epistemic goal (e.g. Alston's) can be achieved in epistemically irrational ways, and even a formulation of the epistemic goal designed to avoid such problems (Foley's) still runs into them. If some of our best attempts at formulating an epistemic goal fail to capture what epistemic rationality is about, then, perhaps epistemic rationality just is not about achieving an epistemic goal.

3.2. A regress argument

The second argument against ICER is adapted from Siegel (1996). Siegel argues against Ronald Giere's and Larry Laudan's instrumental conceptions of epistemic and scientific rationality in particular, but the argument applies to any conception on which epistemic rationality is entirely instrumental in character.

Siegel's question is the following: given means M, evidence E, and goal G, how is it that M can be instrumentally rational as a means to achieve G? The answer is that E must make the following claim rational to believe: "M is an effective means to achieve G" (call this claim 'C'). If E does not make C rational to believe, then M is not rational to adopt as means to achieve G **[iii]**. The mere fact that M will achieve one's goals is not enough to render the adoption of M rational; it must also be rational for one to think that M will do so. The point is perfectly general: for any means M (whether it be a belief or an action) and goal G (be it a practical or an epistemic goal), it cannot be instrumentally rational to adopt M in order to achieve G unless it is epistemically rational for the agent in question to think that M will achieve G.

It is impossible, therefore, for epistemic rationality to be instrumental in character. Instrumental rationality is always about taking the means to achieve our goals, and taking the means to achieve our goals can only be rational if it is epistemically rational to think that the means will achieve our goals. Even in the case where G is an epistemic goal, and M is a belief, it will not be instrumentally rational for a subject S to adopt M unless it is epistemically rational for S to believe that M will achieve G; the mere fact that a belief is instrumentally useful for achieving the epistemic goal is not enough to make it rational to believe.

Another way to put the point is as follows. We can (and sometimes do) hold true beliefs without good reason. We have unfounded hunches, we are wishful thinkers, etc. When such beliefs are true, they serve the epistemic goal, but because we have no good reason to think that they serve the epistemic goal (we have no reason to think that they are true), they are not epistemically rational.

It does not help to object here that if claim C, the claim that M will achieve G, is itself instrumentally effective in achieving the epistemic goal, that is enough to make C epistemically rational, which in turn is enough to make M instrumentally rational. That C is instrumentally effective in achieving the epistemic goal is not enough to make C epistemically rational: what is also required is that there be good reason to think that C is instrumentally effective for achieving the epistemic goal. If there were no good reason to think that, then the case would be analogous to the case where a subject has a true belief that lacks justification. It would be just lucky that the subject's belief C serves the epistemic goal; epistemic rationality would be absent, in that case.

And a good reason to believe C cannot be only a further instrumental reason C*, whose content is that C is instrumentally effective for achieving the epistemic goal. C*, if its own rationality is only instrumental, would depend for its rationality on the further claim C**, whose content is that C* is instrumentally effective for achieving the epistemic goal. And now we're obviously off on an infinite regress of purportedly instrumentally rational beliefs whose rationality depends on the instrumental rationality of higher-level beliefs.

Instrumental rationality, therefore, always depends on the epistemic rationality of the claim that the means are good for achieving the goal, even in the case where the instrument is a belief and the goal is to believe truths. And the epistemic rationality of the claim that the means are good ones for achieving the goal cannot itself be merely instrumental, in the service of the epistemic goal, because that generates a vicious regress.

4. Conclusion.

It is important to see that epistemic rationality is not dependent on the content of the epistemic goals that agents want to achieve. If ICER was correct, then agents could “escape” the dictates of epistemic rationality by lacking an epistemic goal. All that we could say would be that, if an agent cares about achieving the epistemic goal, then that agent has a reason to try to achieve it. Otherwise, nothing is epistemically rational or irrational for her. But that is the wrong result: people who lack epistemic goals can still have epistemically rational or irrational beliefs. Furthermore, if epistemic rationality is the important kind of rationality when it comes to assessing an argument’s premises, and the support that its premises lend to its conclusion, then agents could escape the goodness of arguments that they do not like, and they could escape the badness of their own arguments, merely by lacking an epistemic goal. (Granted, that is easier said than done, but it is at least possible, and it seems to me that it does happen.) But, because epistemic rationality is categorical in nature, simply not caring about achieving an epistemic goal does not allow agents to escape epistemic evaluation of beliefs, or of arguments.

NOTES

[i] For Foley, it is not necessary that a subject actually have considered the argument in question in order for his beliefs to be supported by it. All that is required is that the argument be one that the subject would become aware of, just by reflecting on his reasons for the belief.

[ii] One might object: we should read Foley as holding that each belief must in and of itself satisfy the goal. That might eliminate counterexamples like mine. But that appears to be a major revision of Foley’s account, given that he doesn’t say that anywhere, and given also that the “in and of itself” restriction makes the synchronic restriction pointless. If a subject would take a belief to satisfy the diachronic epistemic goal in and of itself, then he must have an uncontroversial argument for it. So I doubt that Foley had this in mind.

[iii] I want to leave open the question regarding whether the agent must in some sense believe C, or whether it is only the case that she must have grounds that would justify C if she were to form the belief C. Siegel, at least, does not make it clear what he thinks on this point, and I do not want to try to settle the issue one way or the other.

REFERENCES

- Alston, W. (1985). Concepts of epistemic justification. *The Monist*, 68, 57-89.
- David, M. (2001). Truth as the epistemic goal. In M. Steup (Ed.) *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp.151-169). New York: Oxford University Press.
- Feldman, R. (1988). Epistemic obligations. *Philosophical Perspectives*, 2: *Epistemology*, 235-256.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, Mass: Harvard University Press.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66(3), 612-640.
- Kvanvig, J. (2005). Truth is not the primary epistemic goal. In E. Sosa & M. Steup (Eds.) *Contemporary debates in epistemology* (pp.285-296). Oxford: Blackwell.
- Latus, A. (2000). Our epistemic goal. *Australasian Journal of Philosophy*, 78(1), 28-39.
- Nozick, R. (1993). *The Nature of Rationality*. Princeton, NJ: Princeton University Press.
- Siegel, H. (1996). Instrumental rationality and naturalized philosophy of science. *Philosophy of Science*, 63, Supplement: Proceedings of the 1996 Biennial Meetings of the Philosophy of Science Association. Part I: Contributed Papers (Sep., 1996), S116-S124.
- Vahid, H. (2003). Truth and the aim of justification. *Teorema*, 22(3), 83-91.