Published by Blackwell Publishing Ltd, 9600 Garsington Road, Oxford OX4 2DQ, UK, and 350 Main Street, Malden, MA 02148, USA
METAPHILOSOPHY
Vol. 36, No. 4, July 2005
0026-1068

IS CRITICAL THINKING EPISTEMICALLY RESPONSIBLE?

MICHAEL HUEMER

Abstract: There are at least three strategies we might take in approaching controversial issues: (i) we might accept the conclusions of experts on their authority, (ii) we might evaluate the relevant evidence and arguments for ourselves, or (iii) we might give up on finding the answers. Students of "critical thinking" are regularly advised to follow strategy (ii). But strategies (i) and (iii) are usually superior to (ii), from the standpoint of the goal of gaining true beliefs and avoiding false ones.

Keywords: critical thinking, epistemic responsibility, appeal to authority, informal logic, applied epistemology.

1. An Issue for Applied Epistemology[1]

Suppose you have taken an interest in some publicly discussed issue—for instance, the morality of abortion, or the desirability of gun-control legislation, or creationism versus evolution. You would like to acquire a true belief about the issue, or, failing that, at least avoid a false one. That is, you prefer true belief to no belief to false belief. Assume that you know that many experts have published their views on the subject in readily accessible places. These "experts," by definition, are individuals who are intelligent and well informed about the issue and have spent considerable time studying it. Assume also that you are a generally ordinary person, lacking any exceptional cognitive advantages. You do not, for instance, have a genius-level IQ or access to important evidence about the issue that most experts lack. How should you go about deciding what, if anything, to believe?

Consider three strategies:

Credulity: You canvass the opinions of a number of experts and adopt the belief held by most of them. In the best case, you find a poll of the experts; failing that, you may look through several books and articles and identify their overall conclusions.

[1] Battersby (1989) has introduced the notion of applied epistemology, arguing that critical thinking should be identified with applied epistemology.

> *Skepticism:* You form no opinion, that is, you withhold judgment about the issue.
>
> *Critical Thinking:* You gather the arguments and evidence that are available on the issue, from all sides, and assess them for yourself. You try thereby to form some overall impression on the issue. If you form such an impression, you base your belief on that. Otherwise, you suspend judgment.

Granted, there is more to "critical thinking" as traditionally understood than the above description suggests. I assume only that critical thinking requires *at least* this much: that one attempt to assess arguments and evidence on their merits, as opposed to relying on the intellectual authority of others. Thus, critical thinkers do not rely on the fact that a given expert thinks the answer to the issue is so-and-so; critical thinkers look only to the *reasons* the expert has for giving that answer. If they find those reasons wanting, then the expert's opinion will carry no weight with them. And even if they find the reasons cogent, the fact that the reasons were endorsed by the expert will give no *additional* force to the conclusion that they support.

The received view is that Critical Thinking is the best of the three strategies, and certainly better than Credulity. Here are some samples, taken from introductory textbooks, of the kind of advice professors commonly give students:

> [T]he aim of this book is not to offer solutions to a set of ethical dilemmas, but to encourage readers to do the thinking for themselves about these issues. (Thomson 1999, 2)
>
> In this conversation, all sides of an issue should receive a fair hearing, and then you, the reader, should make up your own minds on the issue. (Pojman 1991, 5)
>
> My hope is that exposure to this argumentative give-and-take will encourage students to take part in the process themselves, and through this practice to develop their powers of philosophical reasoning. (Feinberg 1996, xi)

Students should think difficult issues through for themselves—many will consider this a platitude. But there is at least a prima facie puzzle here. Suppose a friend of yours has recently developed chronic abdominal pains. He asks for your advice. You say: "Don't just take the word of some doctors. Diagnose and decide how to treat the condition yourself." Few would consider this to be good advice from the standpoint of maintaining or improving your friend's health. We recognize that there are experts who are better positioned than we are to determine the correct treatment for a medical condition, and we accept the rationality of

deferring to experts about medical issues. What is different in the case of controversial, publicly discussed issues?

I suggest that there may be no relevant difference here—that Critical Thinking, in the kinds of cases I have described, may be unwise in much the same way that diagnosing one's own illnesses is unwise.

2. How Reliable Is Critical Thinking?

Suppose you want to determine whether *a* is *F*, and you have two *F*-measuring devices. One is 80 percent accurate (80 percent of the time it correctly reports whether a thing is *F*). The other is 60 percent accurate. Suppose you cannot use both. Then you would presumably prefer to use the 80 percent reliable device. Alternately, you might suspend judgment on whether *a* is *F*, if you consider an 80 percent probability to be insufficient for adopting a belief. In no event would you prefer the 60 percent reliable device to the 80 percent reliable one. Nor are matters changed if one of the "measuring devices" is one of your natural faculties, such as your eyes. Suppose you can sometimes tell whether objects are *F* just by looking at them. But you also have an *F*-measuring instrument that is more reliable than you are. Then, if you have to choose between using your eyes and using the instrument, you should use the instrument.

Similarly, in choosing between Credulity and Critical Thinking, we will want to consider which approach is more likely to yield true beliefs and avoid false ones. Suppose, first, that the issue in question is one about which there is a consensus among the experts (for example, the issue of evolution versus creationism). Then you should adopt Credulity, accepting that consensus. If instead you adopt Critical Thinking, one of three things will happen:

a) You come to agree with the consensus of experts. In this case, you gain no advantage over Credulity, from the standpoint of securing truth and avoiding error.
b) You come to disagree with the consensus. It is reasonable to think that, in this case, the experts would nevertheless be correct. By hypothesis, the experts are intelligent and well informed and have devoted considerable time and energy to studying the issue. By hypothesis, you have no exceptional cognitive advantages relative to them. Therefore, it seems that any given expert would be no more likely than you are to be in error; even more clearly, the community of experts as a whole is far more likely to be correct than you are.[2]

[2] This observation is supported by Condorcet's Jury Theorem, according to which the probability of the majority being correct is given by

$$\frac{v^{h-k}}{v^{h-k} + e^{h-k}}$$

c) You end in a state of suspended judgment. In this case, you would probably have forgone the opportunity to gain a true belief. For reasons similar to those given immediately above, it seems more likely that your failure to accept the experts' consensus would be due to a mistake or oversight on your part, rather than to a mistake on the part of the community of experts (provided that the threshold level of evidence at which *you* move from withholding to endorsing a belief is not much higher than that of most experts).

Next, suppose the issue is one about which there is no consensus among the experts (for instance, the desirability of gun control). In this case, Skepticism is advisable. If instead you adopt Critical Thinking, one of two things will happen:

a) You will be unable to form a clear impression of the issue, resulting in a state of suspended judgment. In this case, you gain no advantage relative to adopting Skepticism right at the start.
b) You will form a determinate belief on the issue. Should such a belief be trusted? By hypothesis, the experts, with their cognitive advantages, have been unable to form any consensus. This suggests that typical experts are not reliable with respect to the given issue. As a result, it seems unlikely that you would be reliable on the subject either.

How strong a consensus must exist before one adopts Credulity instead of Skepticism will depend upon how one weights the goal of gaining more true beliefs relative to that of avoiding false beliefs—the stronger one's aversion to error, the more often one should rest with Skepticism. But it is difficult to see how Critical Thinking could turn out, in any case, to be the best of the three proposed strategies.

3. Is the Theory of Critical Thinking Consistent?

Suppose you would like to have a true belief about moral realism but lack the time to study the issue. During your office hours, an undergraduate student informs you that she has recently made a thorough study of the issue and has concluded that ethical noncognitivism is correct. You know that the student is sincere, has successfully completed a critical thinking course at your university, has done her best to think critically about the issue, and is otherwise normal. Would you now accept noncognitivism?

where v is the probability of any given expert being correct, e is the probability of a given expert being wrong, h is the number of experts in the majority, and k is the number in the minority. The theorem assumes that experts form their assessments independently and are choosing between exactly two alternatives. Condorcet notes that when we assume individual experts are just 80 percent reliable and the majority outnumbers the minority by as few as nine persons, the probability of the majority being correct exceeds 99.999 percent. See McLean and Hewitt's introduction in Condorcet 1994, 35–36.

As practitioners of critical thinking, we face a dilemma here. On the one hand, if we answer yes, then it seems that we are not using critical thinking in forming our own beliefs. Moreover, the affirmative answer is implausible—who will say that the reliability of an average undergraduate, even after completing a critical thinking course, is high enough that one can have a reasonable belief in noncognitivism solely on the basis of hearing an undergraduate's endorsement of noncognitivism?

On the other hand, suppose we answer no. This seems to imply a negative assessment of the reliability of the techniques of critical thinking as applied by an average student. If those techniques are reliable, then we should assign a high degree of belief to noncognitivism conditional on the fact that a student applying those techniques has endorsed noncognitivism. But if those techniques are not reliable (in the hands of an average student), then why do we advise students to rely on them?

Well, perhaps *the student* rationally estimates his own reliability to be greater than *you* estimate it to be. But it is hard to see why this might be so—and if it is, shouldn't you try to correct the student's estimate, perhaps by supplying him with more evidence relevant to his reliability? For instance, you may know that many students and philosophers have studied metaethics, applying critical-thinking techniques to the best of their ability but coming to a variety of different opinions, many of which conflict with noncognitivism. This might function for you as a defeater for what would otherwise be a justification for believing the student's assessment to be correct. But it is unclear why the student, if apprised of the same facts, should not be equally moved by this consideration as you are.

The peculiarity of the theory of Critical Thinking is that it posits an agent-centered epistemic norm: it holds that, if a person applies certain techniques in arriving at a conclusion, then *she* has good reason to accept that conclusion, but others who know that she arrived at the conclusion by those techniques do *not* thereby have good reason to accept it. It is unclear why this should be so.

4. Objections

Objection 1

Perhaps you should adopt a compromise between Critical Thinking and Credulity: you weigh the opinions of various experts on the issue *and* you also make your own assessment of the evidence and arguments on which the experts base their opinions. You base your conclusion on some kind of weighting of both your and the experts' impressions. This strategy retains a role both for critical thinking and for reliance on epistemic authority. It seems superior to either approach in isolation, since it takes into account the most relevant information, and more information should improve your epistemic performance.

But notice that in such a procedure your own impression of what the evidence supports should have less weight than the impression of a typical expert, given our arguments above. Thus, there is no benefit in going to the trouble of making your own first-hand review of the evidence, unless and until you have gathered the assessments of all the experts. Even at that point, given the large number of experts and their general cognitive advantages, your own impression would have very little relative weight in the overall assessment of the issue.

Still, your own assessment of the direct evidence couldn't *hurt*, could it? For an ideal Bayesian agent, more information would generally be expected to improve epistemic performance.[3] Not so for ordinary human beings, who are subject to biases, mistakes, and so on. In particular, humans can be tempted unjustifiably to weight their own impressions more heavily than the impressions of others.[4] Consequently, if you attempt to implement the compromise strategy, there is a strong danger that you will attach too much weight to your own impressions, relative to those of the more reliable experts. Since the potential epistemic gain is insignificant, it does not justify this risk.

Just as professional doctors are better qualified than ordinary people to diagnose illnesses, professional intellectuals are better qualified than ordinary people to assess the cogency of complex evidence and arguments. Though it is often wise for a patient to seek a second opinion from another doctor, there is no plausible case for patients' attempting to make their own diagnoses—not even if they *also* take into account the opinions of a few doctors.

Objection 2

Perhaps the purpose of teaching critical thinking is not so much for students to acquire the correct answers as for training future academics and intellectuals. Unless students learn to think critically in their philosophy classes, the next generation of philosophers will either not exist or not be able to do valuable research.

The implied piece of practical reasoning here is open to question. The proportion of students in a typical philosophy class who will or should become professional philosophers is quite small. It therefore seems unreasonable to gear one's teaching strategy toward them. In addition,

[3] Suppose E is relevant to whether H, and you do not yet know whether E. If H is true, say the truth-value of H is 1; if H is false, say the truth-value of H is 0. Assume you are a Bayesian agent, P_i is your present probability distribution, and P_f is the distribution you will have after you have found out whether E. Then the expected value (based on P_i) of the *difference* between the truth-value of H and $P_f(H)$ is less than the expected value (based on P_i) of the difference between the truth-value of H and $P_i(H)$: that is what I mean by saying that additional information would be expected to improve your epistemic performance.

[4] Gilovich (1991) discusses biases and other sources of error we are subject to, including the rather extreme tendency to overrate one's own abilities in various respects (77–78).

given the present facts about supply and demand in the field, the worry about a future shortage of philosophers seems farfetched.

However, this reasoning is in any case immaterial to our present concern. Our issue here is the *epistemic rationality* of critical thinking. That a future benefit eventually results, even a benefit in terms of increased knowledge, from practicing critical thinking would not serve to show that critical thinking is epistemically rational.

Objection 3

Our reasoning so far has assumed that your only relevant goals are those of forming true beliefs and avoiding false ones. Perhaps you have other epistemic values, which might be better served by Critical Thinking than by the alternative strategies. What might they be?

First, perhaps you value acquiring *reasonable* beliefs. But just as the average person has no reason to expect his own beliefs, formed by Critical Thinking, to be more often *true* than those of experts, so he has no reason to expect them to be more often *reasonable* than those of experts. This is because the experts are highly intelligent and have devoted much time and energy to the issue on which they are experts, while the average person has no exceptional cognitive advantages. Moreover, if the preceding arguments of this article have succeeded, then they also show that it is at least as reasonable to form beliefs by Credulity as by Critical Thinking. So the desirability of reasonableness does not support Critical Thinking over Credulity.

Second, perhaps you value *knowledge*. But if the experts are reliable sources of information, then it is unclear why true beliefs formed by reliance on their authority should not count as knowledge. And if the experts are not reliable sources of information, then you are probably not reliable either, and so Critical Thinking will not yield knowledge.

Perhaps knowledge requires *justification*, in addition to reliability. But—provided that you grasp the preceding arguments of this article and thus realize that Credulity is more reliable than Critical Thinking—it is also true that you are *more justified* in accepting beliefs produced by Credulity than in accepting beliefs produced by Critical Thinking. (Compare: if you know that one measuring device is more reliable than another, then, ceteris paribus, you are more justified in believing the results produced by the one than those produced by the other.)

Objection 4

Perhaps the thesis of this article is self-undermining, since it implies that an ordinary person should not attempt to evaluate the thesis itself on the basis of the reasons I have presented. Instead, one should consult the opinions of experts. It appears most experts on the subject—such as

professors of philosophy, particularly of logic and epistemology—currently endorse the method of Critical Thinking over that of Credulity. Therefore, if we apply Credulity to evaluating the thesis of this article, we will conclude that we ought to apply Critical Thinking. A potential paradox lies in the offing—suppose that the reader, in accordance with the preceding reasoning, proceeds to evaluate my argument on its merits and finds it to be convincing. In that case, the Critical Thinking approach would also be self-undermining. But let us leave that issue aside.

The important point is that my argument in this article is addressed not to the ordinary person but to the community of experts—I am proposing an argument to professional philosophers, to the effect that they ought not to advise lay people to think critically about certain kinds of issues. The experts themselves cannot evaluate my thesis on the basis of the opinions of the experts until a sufficient number of them have first evaluated my thesis on its merits. Hence the function of the arguments I have presented. If the experts should be convinced by my argument, they would then go on to expound it to laypeople, who would then be justified in accepting it.

Objection 5

It seems that, even if one wants to rely primarily on the judgments of experts to form beliefs about difficult issues, one must still rely to some extent on critical thinking, in order to identify trustworthy experts to begin with. Thus, one cannot make a blanket rejection of the method of critical thinking.

I agree with this objection, and this points up the need to clarify my thesis in this article. It would be foolish to argue that, in general, one should not rely on one's own judgment and reasoning; indeed, it is questionable whether such advice could in principle be taken. I contend only for the more modest thesis that, with respect to publicly discussed issues, one should usually not rely on one's own judgment and reasoning directly about the publicly discussed issue itself. This is compatible with the point that one *should* rely upon one's own judgment and reasoning in determining which experts to rely on with regard to the publicly discussed issue.

Why does it make more sense to use one's own reasoning to determine who is a qualified expert than to use one's own reasoning to determine, directly, what the answer to the publicly discussed issue is? Why, for instance, is it more reasonable for me to trust my judgment as to whether Professor Kleck is an expert on gun control than to trust my judgment as to whether gun control is beneficial? One answer to this would be that there is no body of experts on the question of who is an expert on gun control, although there is a body of experts on the question of whether

gun control is beneficial. However, the more important answer is that it is much easier to determine whether a particular individual is an expert on gun control than it is to determine whether gun-control laws are beneficial—in other words, the former question does not *require* expertise in the same way the latter does.

Consider an analogy. Courts and legislatures often call expert witnesses to testify. Why do they not also need to call "second-order experts," to determine whether the first experts are really experts? The answer is that, with respect to issues where the notion of "expertise" is relevant, it is generally much easier to determine who are the experts with regard to the issues than it is to determine the answers to the issues directly. Thus, a court can determine whether a particular witness is an expert on a particular medical technology, even when the court itself lacks expertise with regard to that technology. The approach I am recommending in this article is analogous to that taken by courts and legislatures in relying on expert testimony. I simply take a more expansive view than usual with regard to the class of issues calling for expertise, arguing that they include the likes of abortion, gun control, and most philosophical issues.

5. When Is Critical Thinking Rational?

I do not deny that critical thinking is valuable in some situations. If one is considering an important issue that is not publicly discussed—say, a decision one faces in one's personal life—then Credulity is unavailable and Critical Thinking may be required.

Here is another situation in which Critical Thinking may give the nonexpert her best chance of reaching a true conclusion. Suppose you believe that a particular issue is not difficult to resolve if one applies the methods of Critical Thinking but that a substantial proportion of the experts—perhaps owing to bias on their part—have not made a serious effort to apply those methods. In such a case, my central arguments against Critical Thinking would not apply. The intelligence of the experts, the amount of information available to them, and the amount of time and effort they have devoted to the issue may not make up for the disadvantage created by their bias. Indeed, Kornblith (1999) has argued persuasively that intelligence can *exacerbate* the problems created by bias.[5] Given this, one could have grounds for anticipating one's own exercise in Critical Thinking to be more reliable than the experts' assessments.

This approach to defending Critical Thinking is not *merely* an appeal to the possibility of biased experts. Such an appeal would be ineffectual in

[5] Essentially, Kornblith's argument is that intelligence can be used as a tool for helping one to rationalize false but pleasing beliefs.

the absence of evidence that ordinary people tend to be less biased than typical experts. Nor can the idea simply be that many of the experts have failed to practice the methods of critical thinking through *inadvertent error*. This hypothesis would not avail the defender of critical thinking unless it could be argued that ordinary people are less prone to inadvertent errors of this kind.

It seems, then, that the sort of consideration suggested favors the adoption of Critical Thinking only if (a) something about the experts renders them less able than ordinary people to implement the techniques of critical thinking, or (b) the experts have not generally tried to implement those techniques. There may well be cases in which one or the other of these conditions holds. If there are, and the layperson has good reason to believe he is dealing with such a case, then the approach of Critical Thinking is probably his best bet.[6]

Department of Philosophy
University of Colorado
Boulder, CO 80309–0232
USA
owl232@earthlink.net

References

Battersby, Mark. 1989. "Critical Thinking as Applied Epistemology: Relocating Critical Thinking in the Philosophical Landscape." *Informal Logic* 11:91–100.

Condorcet, Marquis de. *Condorcet: Foundations of Social Choice and Political Theory*. Edited and translated by Iain McLean and Fiona Hewitt. Aldershot, U.K.: Edward Elgar Publishing.

Feinberg, Joel, ed. 1996. *Reason and Responsibility: Readings in Some Basic Problems of Philosophy*. 9th ed. Belmont, Calif.: Wadsworth.

Gilovich, Thomas. 1991. *How We Know What Isn't So*. New York: Free Press.

Kornblith, Hilary. 1999. "Distrusting Reason." *Midwest Studies in Philosophy* 23:181–96.

Pojman, Louis, ed. 1991. *Introduction to Philosophy: Classical and Contemporary Readings*. Belmont, Calif.: Wadsworth.

Thomson, Anne. 1999. *Critical Reasoning in Ethics*. London: Routledge.

[6] I would like to thank Josh Snyder, Sheralee Brindell, Ben Kovitz, an anonymous referee for *Metaphilosophy*, and my numerous students who have heard the argument of this article, for their comments and criticisms.