David Coady

When Experts Disagree

ABSTRACT

Alvin Goldman has criticized the idea that, when evaluating the opinions of experts who disagree, a novice should "go by the numbers". Although Goldman is right that this is often a bad idea, his argument involves an appeal to a principle, which I call the *non-independence principle*, which is not in general true. Goldman's formal argument for this principle depends on an illegitimate assumption, and the examples he uses to make it seem intuitively plausible are not convincing. The failure of this principle has significant implications, not only for the issue Goldman is directly addressing, but also for the epistemology of rumors, and for our understanding of the value of epistemic independence. I conclude by using the economics literature on information cascades to highlight an important truth which Goldman's principle gestures toward, and by mounting a qualified defense of the practice of going by the numbers.

> As if someone were to buy several copies of the morning paper to assure himself that what it said was true. (Ludwig Wittgenstein, *Philosophical Investigations*, 265.)

1. EXPERTS AND NON-DISCRIMINATING REFLECTORS

In the course of a discussion of our epistemic dependence on experts, Alvin Goldman (2002, 151) objects to the procedure, advocated by Lehrer and Wagner (1981, 20), of seeking rational consensus by identifying a "weighted average" of worthwhile opinions, where a worthwhile opinion is defined as one that is better than a random device. Although Goldman is right to object to Lehrer and Wagner's procedure, his argument against it appeals to a principle – which I call the non-independence principle – which is not in general true. An analysis of the limitations of the non-independence principle has significant implications, not only for how we should understand our epistemic dependence on experts, but also, more generally, for how we should understand our epistemic dependence on others. Not everyone on whom we are epistemically dependent is (in any ordinary sense) an expert.

What should a novice or layperson, that is a nonexpert, think in a situation in which experts disagree?[1] Goldman (2002, 150) notes that it is tempting to appeal to the numbers on either side of the issue:

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

> Each new testifier or opinion-holder on one side of the issue should add weight to that side. So a novice who is otherwise in the dark about the reliability of the various opinion-holders would seem driven to agree with the more numerous body of experts. Is that right?

According to Goldman, it is not. As we shall see, Goldman is right to object to this position, but he is right for the wrong reason. His reason is the non-independence principle:

> If two or more opinion-holders are totally *non-independent* of one another, and if the subject knows or is justified in believing this, then the subject's opinion should not be swayed – even a little – by more than one of these opinion-holders. (Goldman 2002, 151)

The concept of non-independence is expressed in terms of conditional probability. Suppose there are two opinion-holders X and Y and an hypothesis H, and let X(H) be X's believing H and Y(H) be Y's believing H. To say that Y's belief is totally non-independent of X's belief is to say:

$$P(Y(H)/X(H)\&H) = P(Y(H)/X(H)\&\sim H)$$

In other words, Y would be just as likely to share X's opinion that H when it is false as she is when it is true. In this situation, Goldman calls Y a *non-discriminating reflector* of X with respect to H. According to Goldman, in order for Y's opinion to have any evidential worth above and beyond X's opinion it is necessary for Y to be more likely to share X's opinion that H is true when it is true than when it is false. In other words:

$$P(Y(H)/X(H)\&H) > P(Y(H)/X(H)\&\sim H)$$

In this situation Y's belief is at least partly *conditionally independent* of X's belief.

The non-independence principle may seem plausible. Consider an extreme case, which Goldman discusses (2002, 151-2). Suppose X is a 'guru' and Y is a 'follower' who believes whatever X believes. If X believes H, Y is certain to believe H. Hence Y is just as likely (that is, with probability 1) to share X's opinion if it is true as he or she is if it is false:

> ... a follower's opinion does not provide any additional grounds for accepting the guru's view (and a second follower does not provide any additional grounds for accepting a first follower's view) even if all the followers are precisely as reliable as the guru himself (or as one another) – which followers must be, of course, if they believe exactly the same things as the guru (and one another) on the topics in question.

Goldman concludes (2002, 154) that Y's agreement with X about H would only provide evidence of H for a third person if that person has evidence that Y has a "more-or-less autonomous causal route to belief, rather than a causal route that guarantees agreement with X". He mentions three forms such autonomy could take - access to independent eyewitnesses, access to independent experiments, and a process of reasoning with X about the truth of H. The presence of some autonomy in any of these forms would make Y "poised to avoid belief in H even though X believes it".

People who value epistemic autonomy (as I hope we all do) are likely to associate the language of "gurus" and "followers" with connotations of irrationality which may prejudice clear discussion of the issue. It is important to remember that on Goldman's

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

formal account one is only a non-discriminating reflector with respect to *a particular belief* of another. You could be a non-discriminating reflector of another person with respect to one of that person's beliefs and still be as discriminating as you like about any or all of his or her other beliefs.

It is also important to note that Goldman quite rightly does not think that there is anything wrong, as such, with being a non-discriminating reflector. In fact, the problem he is addressing, about which putative experts a novice should believe, is a problem about which putative experts a novice should non-discriminately reflect. Goldman stipulates (2002, 144) that the novice has no prior opinions (or at least none to which he feels he can legitimately appeal) about the domain in which the putative experts claim expertise. Hence the novice can have no autonomous causal route (in Goldman's sense) to the belief he or she ends up with.

An illustration may be helpful here. Suppose that I have just travelled to a new community and am told by X, who purports to be a meteorologist, that it will be hot tomorrow (proposition H). Suppose, as a result, I attach some credibility to H. Next I encounter Y who also believes that it will be hot tomorrow. Goldman's position implies that I would be irrational to allow Y's belief to increase my confidence that it will be hot tomorrow, if Y is a non-discriminating reflector of X with respect to H. That is, according to Goldman, I should only take Y's belief as confirming evidence that H, if Y's belief that H is at least partly conditionally independent of X's belief that H. Furthermore, the same would be true no matter how many people I find in this community who believe that tomorrow will be hot. As Goldman says, it makes no difference how many people share an initial expert's opinion: "If they are all non-discriminating reflectors of someone whose opinion has already been taken into account, they add no further weight to the novice's evidence" (2002, 154).

Goldman presents his formal argument for this position in Bayesian terms. The novice should update his or her belief in H in the light of the evidence that X believes H in accordance with the following "likelihood quotient":

$$(1)\quad \frac{P(X(H)/H)}{P(X(H)/\sim H)}$$

And the novice should update his belief in H in the light of the evidence that X and Y believe it in accordance with this likelihood quotient:

$$(2)\quad \frac{P(X(H)\&Y(H)/H)}{P(X(H)\&Y(H)/\sim H)}$$

Now, normally you would expect (2) to be larger than (1), but as Goldman spills some ink demonstrating, when Y is a non-discriminating reflector of X with respect to H, that will not be so: (1) and (2) will be equivalent.

Although mathematically impeccable, this argument involves a questionable assumption, namely that the probabilities in question, and hence the ratios between them, will remain constant. Perhaps the clearest way to see why this cannot be assumed is to concentrate, not on the updating process itself, but on its end result, the degree to which the novice should believe H (that is, the probability the novice should assign H).

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

After learning that X believes H, the novice updates his or her degree of belief in H by assigning it the 'new' probability P(H/X(H)). After learning that Y believes H as well, the novice again updates his or her degree of belief in H; this time assigning it the probability P(H/X(H)&Y(H)). In general, the latter value will be greater than the former; however, if Y is a non-discriminating reflector of X with respect to H, these values will be the same. Hence, in this case, 'updating' the novice's degree of belief from the former to the latter should not increase his or her confidence in H.

The presupposition of this argument that P(H/X(H)) will remain constant for the novice throughout the enquiry is not justified, because Bayesian probabilities are subjective. They are measurements of the degree to which it is rational to believe something, given certain evidence. Hence, probabilities can change. It is therefore wrong to think of a Bayesian probability as something that has, in Goldman's words, "already been taken into account". Goldman's argument assumes, in effect, that knowing or justifiably believing that P(H/X(H)) and P(H/X(H)&Y(H)) are equivalent is a reason for the novice to assign a lower value than he or she otherwise would to the latter. But this ignores the possibility that it might instead be a reason to assign a higher value than he or she had previously assigned to the former.[2]

Why should we consider the latter possibility? Because the existence of a non-discriminating reflector of a person with respect to a proposition can itself be evidence in favour of that proposition. Suppose that Y is a non-discriminating reflector of X with respect to H, *because Y knows or is justified in believing that H is within a domain in which X is an expert*. Y believes H because X does, and would do so even if H were false, but Y's concurrence with X still provides the novice with evidence for H, because the novice rationally believes Y to be a reliable judge about whether X is a reliable judge about whether H is true. The novice's confidence in X's expertise concerning H is rationally increased by his or her confidence in Y's meta-expertise.[3] Y's meta-expertise consists in Y's knowledge of (or justified belief about) the scope and extent of X's expertise.[4] Hence, contrary to the non-independence principle, it may well be rational for a subject, who knows or is justified in believing that two or more opinion holders are totally non-independent of one another, to be swayed - perhaps quite a lot - by more than one of them.

I anticipate the following objection to my argument, which concedes its validity, but questions its significance.[5] It is true that in the above circumstances a novice should have his or her confidence in H increased by Y. But once the fact that Y is totally non-independent of X has been taken into account, the further fact that Y actually believes H should not increase the novice's confidence in H. This suggests that a fairly simple alteration of the non-independence principle might avoid the preceding objection.

But even if that's true, the epistemic situation facing novices is not, in general, divisible in a way that would make such a reformulated principle very useful. Suppose, to return to my earlier example, that X is, as he claims, a meteorologist. Suppose further that everyone else in the community believes what he says about tomorrow's weather because they know him to be well qualified and invariably accurate in his weather predictions up until now. My confidence that it will be hot tomorrow may be rationally increased by the fact that many apparently sensible people believe it. In making this assessment I may quite

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

rationally be indifferent to the issue of how (conditionally) independent, with respect to that belief, they are from one another.

In fact, the degree by which my confidence is increased by the concurrence of the many may be greater if I know that all but one of them are non-discriminating reflectors of the other, than it would have been if I had known instead that they all had a partly autonomous causal route to their belief. Suppose, for example, that the only even partly autonomous causal routes to belief available to the non-discriminating reflectors are intuitive inductions from their own personal experience. They may be poor meteorologists, but good judges of meteorologists. And a novice may rationally judge that this is so. Hence, the claim that "a follower's opinion does not provide any additional grounds for accepting the guru's view", is not in general true. It would only be true, if we could presuppose that followers are invariably unreliable judges of gurus. And we can't.

2. RUMORS

Goldman (2002, 151) extends his argument beyond situations in which novices are considering the claims of rival experts:

> Another example, which also challenges the probity of greater numbers, is the example of rumors. Rumors are stories that are widely circulated and accepted though few of the believers have access to the rumored facts. If someone hears a rumor from one source, is that source's credibility enhanced when the same rumor is repeated by a second, third or fourth source? Presumably not ...[6]

There is a widespread view that although there can (and very likely will) be a diminution in the reliability of a communication (whether it is strictly a rumor or not) as it passes from person to person, there cannot be any increase in its reliability. It is certainly an appealing thought, and has been endorsed by numerous philosophers, including John Locke (*Essay*, bk. iv, ch. xvi, s. 10):

> ... *any testimony, the further off it is from the original truth, the less force and proof it has*. The being and existence of the thing itself, is what I call the original truth. A credible man vouching his knowledge of it is a good proof; but if another equally credible do witness it from his report, the testimony is weaker: and a third that attests the hearsay of an hearsay is yet less considerable. So that in traditional truths, each remove weakens the force of the proof: and the more hands the tradition has successively passed through, the less strength and evidence does it receive from them.[7]

It should be clear by now what is wrong with this reasoning. It ignores the fact that each person in the chain along which a communication passes can decide not to pass it on, on the grounds that they don't believe it. What is more, the judgement about whether to believe it may be based on sound considerations about the credibility of a particular informant. To the extent that such credibility judgements operate as selection pressures which contribute to the survival and spread of rumors, the more a rumor is repeated the more likely it is to be true.[8]

Of course, there may be selection pressures other than credibility judgements at work, such as *interest*, and it is possible for rumors to become less credible, or at least for them

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

not to become more credible, as a result of being repeated. But there is certainly no *a priori* reason to believe that repetition cannot cause rumors to become more credible.[9] Indeed there is empirical evidence that, in at least some circumstances, the more often a rumor is repeated the more likely it is to be true.[10] Theodore Caplow (1947, 301), who was given the task of studying rumors among Allied troops during the World War II, reported "a positive and unmistakeable relation between the survival of a rumor, in terms of both time and diffusion, and its veracity".

3. INFORMATION CASCADES

The situations we have been looking at are all examples of what economists call 'information cascades'. The literature on information cascades casts light on the non-independence principle, revealing an element of truth to Goldman's analysis, as well as highlighting its limitations.

An information cascade can occur when people express their beliefs about the answer to a certain question in a sequence. If the early beliefs show a clear pattern, the information inferred from this pattern may outweigh any private information that individuals later in the sequence have which conflicts with it. Hence, they "follow the crowd", rather than follow their own "private" evidence. Information cascades are ubiquitous. On a reasonably fine day I am wondering whether to take an umbrella to work. I look out the window to see if others are carrying umbrellas. If enough of them are, I do too, even though I may reasonably conclude that many of them are only carrying umbrellas because they have seen that others are carrying umbrellas, and even though whatever private information I have indicates that rain is unlikely.

Information cascades have been studied in laboratory conditions. The following experiment (Anderson and Holt 1997) provides a useful context for discussion. At the beginning of the experiment, there are two urns. One of them, the predominantly white urn, contains twice as many white marbles as dark marbles. The other, the predominantly dark urn, contains twice as many dark marbles as white marbles. The two urns are outwardly indistinguishable. One of them is chosen at random, and volunteers are asked to draw one marble from it, in sequence, and predict which urn they are drawing from. Those who predict correctly are rewarded with two dollars. The volunteers do not have any direct information about the colour of marbles drawn earlier in the sequence. They only have two pieces of relevant evidence on which to base their prediction: the colour of the marble that they themselves have drawn, and the predictions of those earlier than them in the sequence.

Now, what should you do when your two pieces of information conflict? Suppose, for example, that you have drawn a white marble, but everyone before you has predicted that the urn is predominantly dark. If there is only one person ahead of you, and you can assume that he or she is rational, then you can infer that he or she has drawn a dark marble. Hence one dark marble and one white marble have been drawn and it is equally likely to be either urn. Most subjects in this situation prefer to rely on their private information, that is, they will predict that it is the predominantly white urn.[11] However, suppose that there are two people ahead of you and they have both predicted that it is the predominantly

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

dark urn. Now, it seems clear that you should agree with the people ahead of you, even if your private information indicates otherwise. You have good reason to believe that the first two marbles are dark, and that outweighs any information you can obtain from the one marble you have drawn. In general, whenever the first two predictions match, the third person should follow, regardless of the colour of the marble he or she draws. This is how an information cascade develops. Not only should the third person follow, so should the fourth, the fifth, and so on, even if their private information indicates otherwise. To use Goldman's language, they will all be non-discriminating reflectors of the second person, with respect to his or her belief that the urn is predominantly dark.

In this situation, it seems that Goldman is right. On the face of it, it would be a mistake for a person with no private information (a novice) to construe the agreement of large numbers of people that the urn is predominantly dark as correspondingly weighty evidence that the urn is in fact predominantly dark. In this situation, the novice's evidence that the urn is predominantly dark consists, it seems, only in his or her evidence that the first two marbles that were drawn from it are dark, which in turn consists only in the fact that the people who drew those marbles believe that the urn is predominantly dark. This evidence is not strengthened by the subsequent agreement of everyone else. In this situation "going by the numbers", to use Goldman's phrase, would lead to overconfidence. It really would be like buying several copies of the morning paper to assure oneself that what it said was true. Or so it would seem.

In another laboratory experiment, devised by Angela Hung and Charles Plott (2001), everything remained the same except for the rules governing how subjects were paid. Instead of being paid if and only if they got the correct answer, subjects were paid if and only if the majority of them got the correct answer. The result was that subjects' predictions were much more likely to be determined by the colour of the marble they had drawn than by the predictions of people earlier than them in the sequence. This was clearly rational. It prevented the development of an information cascade that would conceal private information from those later in the sequence. Hence, by taking, in Goldman's words, "an autonomous causal route to belief," a person can increase the likelihood that the majority is right at the same time as reducing the likelihood that he or she is right.

The work of Hung and Plott has been used to argue that effective collective decision-making requires epistemic autonomy (Surowiecki, 65). One way of achieving this is to get individuals to express their beliefs simultaneously:

> Organizations ... clearly can and should have people offer their judgments simultaneously, rather than one after the other. On a deeper level, the success of the Hung and Plott experiment – which effectively forced people in the group to make themselves independent – underscores the value and difficulty of autonomy. One key to successful group decisions is getting people to pay much less attention to what everyone else is saying.

But this is a mistake, like the one I've argued Goldman makes. Epistemic autonomy in the Hung and Plott experiment increases the value of "going by the numbers" because of some special characteristics of the experimental situation. These characteristics are rarely found in the real world.

The most salient of these characteristics is that everyone has exactly the same amount

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

of private information. In the case I considered earlier of a community with one meteorologist and many non-discriminating reflectors of his belief about tomorrow's weather, the majority would clearly be less likely to be right if they had ignored what others were saying, and paid attention only to their own extremely limited private information. Hence, it would clearly be a mistake for an outsider, who was unable to reliably determine who had meteorological expertise, to get these people to offer their weather forecasts independently, by, for example, getting them to make them simultaneously.

Hung and Plott (2001) claimed that their experiments showed that outside observers learn more when members of a group make their decisions independently. But this only holds in extremely simple and largely artificial situations. Those are situations in which no one in the group is an expert, in the sense that no one in the group has either more private information than others, or a greater capacity than others to make good inferences from private information when answering the question at issue.

4. SOME APPLICATIONS

Remember that Goldman's discussion occurs in the context of a broader discussion of what novices should believe in a situation in which experts disagree about some proposition H. After emphasising again that numbers on either side of the issue need not be decisive (2002, 155), Goldman concludes:

> The appropriate change in the novice's belief in H should be based on two sets of concurring opinions (one in favor of H and one against it), and it should depend on *how reliable* the members of each set are and on *how (conditionally) independent* of one another they are.

He claims (2002, 155-6) that this conclusion seems to get the right result in the following case:

> If scientific creationists are more numerous than evolutionary scientists, that would not incline me to say that a novice is warranted in putting more credence in the views of the former than in the views of the latter (on the core issues on which they disagree). At least I am not so inclined on the assumption that the novice has roughly comparable information as most philosophers currently have about the methods of belief formation by evolutionists and creationists respectively.

In an endnote attached to this passage (2002, 163 n. 21), Goldman makes it explicit that he is specifically "assuming that believers in creation science have greater (conditional) dependence on the opinion leaders of their general viewpoint than do believers in evolutionary theory".

Two things should be said at this point. In the first place, it is not obvious that Goldman's assumption is correct. I, for example, am a believer in evolutionary theory on the core issues on which it contradicts creationism, but my beliefs in this area are all highly conditionally dependent on the opinion leaders of that general viewpoint. I suspect this is true of many other believers in evolutionary theory. But even if Goldman's assumption is correct, our previous discussion makes it clear that this is not *on its own* a reason for a novice to prefer evolutionary theory to creationism. It is only a reason,

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

when combined with the further assumption (which is surely legitimate in this case) that creationist "followers" do not have good reason to believe that creationist "opinion leaders" have exceptionally good access to private information or exceptionally good abilities to make appropriate inferences from that information.

In fact, I don't think this makes a very good case study. Hardly any reasonably educated person can be a complete novice in this dispute, and even people who are relatively uninformed don't have to look far to see internal inconsistencies in the creationist position. A better example is global warming. Most meteorologists believe that it is caused by human activity, a small minority disagree. As things stand many people, including myself, have little but this bare fact to go on when deciding what to believe; nonetheless, I think we are justified in agreeing with the larger group of experts, just because it is larger.[12] Goldman's position implies that if we were to discover that the beliefs of some members of the larger group of experts were highly (conditionally) dependent on others, our confidence in the proposition about which they agree should be reduced. But, as I hope I've made clear, this need not be the case. We may be quite rationally indifferent to the discovery of these dependence relations. We may even quite rationally see them as evidence justifying an increase in our confidence in the proposition in question. It all depends on the details.

Of course, my defence of "going by the numbers" must be qualified. Most of us could do a great deal more than just go by the numbers when choosing which group of experts to trust on the topic of global warming, and Goldman's article is full of valuable insights on how to make choices of this kind. So, I cannot at the moment rule out the possibility that further investigation could make it rational for me to side with the minority on this issue. However, all of the other procedures available to novices trying to determine which experts to trust (for example, examining their qualifications, interests and prejudices, or the quality of their arguments) have the disadvantage of requiring the novice to acquire a degree of expertise or meta-expertise himself. The procedure of "going by the numbers" may well be the only available rational procedure for a novice who lacks the requisite time, ability, or energy to do this. Of course a novice in this position could simply decide to suspend judgement on the topic. But, as the example of global warming makes clear, this can be a rash approach if the issue is of vital importance, and if the novice *must* make important choices, such as how to vote, based on what he or she believes.

I have not explicitly defined what an expert is, or addressed the issue of how a novice is to identify genuine experts. This is partly because I am quite happy with Goldman's definition, and partly because it does not matter much, since my principal concern is with the broader issue of our epistemic dependence on others.[13] But one lesson specifically about expertise can be drawn from the preceding discussion. Someone does not fail to be an expert in a particular domain just because his or her views in that domain are largely dependent on others. Although it seems plausible that genuine experts in the natural sciences, or pure mathematics have a significant degree of independence from one another in their belief forming practices, experts on questions of history, and perhaps more generally, questions in the social sciences, are inevitably highly dependent on their sources.

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

5. Conclusion

So, where does this leave the debate between Goldman on the one hand and Lehrer and Wagner on the other? In what circumstances, that is, do new opinion-holders on one side of an issue "add weight" to that side? Both parties recognise that if the new opinion-holders are utterly unreliable, they add no weight. Unlike Lehrer and Wagner, however, Goldman recognises that there are circumstances in which new opinion-holders can fail to add weight to their side and in which their reliability is not the issue. This is made clear by the phenomenon of information cascades. In an information cascade it seems that an outside observer should not view the agreement of more and more people, as the accumulation of more and more evidence for the proposition about which they agree, even though we may suppose that they are all 'reliable', that is, their opinions are all more likely to be right than a random device would be.

Unfortunately, Goldman's account of the circumstances in which new opinion-holders do not add weight to their side of an issue is far too broad. It may well be rational to be influenced by new opinion-holders, even when their opinion is entirely dependent on someone whose opinion we have "already taken into account". Indeed we may rationally judge it to be decisive.

It may seem that my position implies a devaluing of epistemic autonomy. But I don't believe that's true. On the contrary, my discussion of information cascades gives us a clearer picture of the nature of its value and the circumstances in which it is valuable. The literature on information cascades reports a reasonably widespread tendency on the part of many subjects to ignore the crowd and be guided instead by their own private evidence, even when that is clearly individually irrational. But, as we've seen this kind of individually irrational belief formation can be collectively beneficial. So, the epistemic virtue of what might be called independent thinking in an information cascade can be compared to the moral virtue of cooperation in prisoner's dilemmas. In both cases everyone involved would be better off if everyone were to behave in a way that is individually irrational (at least on one plausible construal of individual rationality) than they would if everyone did not. But both virtues are conditional. We don't want Mafia dons cooperating by keeping their codes of silence or corporations cooperating in engaging in uncompetitive practices, and we don't want people being epistemically autonomous when they could make their views dependent on others who either have much more information or a much greater ability to make rational inferences from their information.

Acknowledgements

Thanks to Alvin I. Goldman, Richard Corry and James Chase for helpful feedback on earlier versions of this paper.

References

Anderson, Lisa R., and Holt, Charles, A, (1997). "Information Cascades in the Laboratory". *American Economic Review* 87: 847–62.

Caplow, Theodore (1947). "Rumors in War". *Social Forces* 25: 298–302.

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

Coady, David (2006). "Rumour has it". *International Journal of Applied Philosophy* 20(1): 41–53.

Goldman, Alvin I. (2001). "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63: 85–109.

—— (2002). "Experts: Which Ones Should You Trust?" In Alvin I. Goldman, *Pathways to Knowledge: Private and Public.* New York: Oxford University Press.

Hung, Angela, and Plott, Charles (2001). "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." *American Economic Review* 91: 847–62.

Laplace, Pierre (1951). *A Philosophical Essay on Probabilities*, trans. F.W. Truscott and F.E. Emory. New York: Dover.

Lehrer, Keith and Wagner, Carl (1981). *Rational Consensus in Science and Society: a Philosophical and Mathematical Study*. Boston, MA: Kluwer.

Locke, John (1947). *Essay Concerning Human Understanding.* London: Dent.

Surowiecki, James (2004). *The Wisdom of Crowds*. London: Little, Brown.

Wittgenstein, Ludwig (1968). *Philosophical Investigations,* trans. G.E.M. Anscombe, Oxford: Blackwell.

NOTES

[1] I shall follow Goldman's practice of referring, for the sake of brevity, to disagreements between experts, when the real issue is disagreements between *putative* experts. The novice may not be in a position to judge whether a putative expert really is an expert. See Goldman 2002, 161 n. 9.

[2] In this argument I've been assuming that the novice discovers that X believes H before discovering that Y believes H. What if the novice simultaneously discovers that X and Y believe H? Just as Bayesian probabilities can change over time, they can vary across possible worlds. So, we cannot assume that knowing or justifiably believing that P(H/X(H)) and P(H/X(H)&Y(H)) are equivalent is a reason to assign a lower value than would otherwise have been assigned to the latter, rather than a reason to assign a higher value than would otherwise have been assigned to the former.

[3] I borrow the term "meta-expert" from Goldman (2002).

[4] Y knows the scope of X's expertise to the extent that he or she knows which questions X can answer with a high probability of accuracy. Y knows the extent of X's expertise to the extent that Y knows how high that probability is.

[5] This was suggested to me by a conversation with Goldman, but I cannot be certain that he intended to make this objection.

[6] I have omitted the end of the last sentence of this quotation, which says "especially if the hearer knows (or justifiably believes) that these sources are all uncritical recipients of the same rumor," because it is the nature of being an uncritical recipient that is at issue here. If being an uncritical recipient of a rumor means being utterly unreliable (*i.e.*, no better than a random device), then the point will be granted, but we should not assume that those from whom we hear rumors are utterly unreliable. If, by contrast, being an uncritical recipient of a rumor instead means being a non-discriminating reflector of the person from whom one heard the rumor *with respect to the rumor's content*, then my previous arguments applies.

[7] Similar thoughts are expressed by Pierre Laplace (Ch. XI, 124).

[8] In academic terms, it has passed an increasingly rigorous process of peer review.

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press

[9] For a discussion of the circumstances in which rumors become more credible the further they spread see (Coady 2006).

[10] When I say it is more likely to be true the further it disperses, I mean that it is more likely that it was true all along. Rumors can also *become true* in one of two ways. Rumors can change to conform to a pre-existing reality. This possibility is suggested by the analogy with natural selection working on random variations, with credibility acting as a possible selection pressure. But a rumor can also become true, because reality changes to conform to the rumor. In other words, some rumors, such as rumors about stock-market collapses or social conflict, can be self-fulfilling.

[11] Arguably this is the rational thing to do, since it is always possible that the person in front of you has made a mistake, and consequently you should have less confidence that the first marble was black than you have that the second marble is white.

[12] I am ashamed to admit how deeply ignorant of this important subject I am, but my ignorance does allow me to illustrate my argument quite effectively.

[13] I believe this is Goldman's principal concern too.

David Coady is lecturer in philosophy at the University of Tasmania, Australia. His research interests include epistemology, metaphysics and political philosophy. He recently edited the book *Conspiracy Theories: The Philosophical Debate*.

https://doi.org/10.3366/epi.2006.3.1-2.68 Published online by Cambridge University Press