Epistemic Free Riders and Reasons to Trust Testimony[1]

Abstract:

Sinan Dogramaci has recently developed a view according to which the function of
epistemic evaluations – like calling someone's behavior 'rational' or 'irrational' – is to
encourage or discourage the behavior evaluated. This view promises to explain the
rational authority of testimony, by describing a social practice that promotes the
coordination of epistemic procedures across a community. We argue that Dogramaci's
view is unsatisfactory, for two reasons. First, the social practice at its heart is vulnerable
to free riders. Second, even if the problem of free riders can be solved, it 'alienates'
epistemic agents from the testimony that they receive, in that, though they <u>will</u> accept
testimony from their fellows, they will have no <u>reason</u> to do so. We argue that a more
satisfactory view can be had if we couple the genuine insights that are to be found in
Dogramaci's proposal with the recognition that testimony is an excludible good that is
often distributed according to market forces. It is this fact about testimony that
discourages free riding, and, when coupled with an amended version of Dogramaci's
view, provides agents with a reason to trust testimony.

1. Introduction

In "Reverse Engineering Epistemic Evaluations" Sinan Dogramaci argues for a

view that he calls 'epistemic communism'. This view has four central components:

(1) Epistemic evaluations encourage or discourage certain behaviors. In particular,

evaluating someone positively (by, for example, calling them 'rational')

encourages that person to continue to use the epistemic procedure they were

employing when they were evaluated, and encourages others to use that procedure

as well. Negative epistemic evaluations tend to have the opposite effect.

(2) 'there is a robust correlation between the rules a speaker (implicitly) evaluates positively/negatively and the rules that she herself does/doesn't follow'[2]

(3) Together (1) and (2) promote coordination of epistemic procedures among community members. Moreover, community members can easily tell that they and their fellows employ the same epistemic procedures.[3]

(4) Together (1), (2), and (3) provide epistemic agents with a reason to trust testimony.

Dogramaci's idea is that epistemic agents will evaluate the epistemic procedures of their fellows, encouraging those of which they approve, and discouraging those of which they disapprove. In the long run, this process encourages a community to employ a single set of epistemic procedures. Any community member will, then, have reason to trust the testimony offered by his or her fellows, since that testimony will reflect an epistemic procedure that he or she employs his/herself.

We believe that there is a lot to like about Dogramaci's proposal, but that, despite its considerable merits, it requires refinement and emendation. As presently formulated, it faces two problems. First, the cooperative practice at its heart is unstable and vulnerable to free riders. The benefits generated by epistemic communism accrue to community members regardless of whether or not they evaluate the epistemic procedures of their fellows. However, there is a cost associated with identifying and evaluating the epistemic procedures that other people use. It is, therefore, rational for each agent to

accept the benefits of epistemic communism without paying its costs. In order to prevent free riding there must be some mechanism that motivates agents to evaluate the epistemic practices of others.

Second, epistemic communism ensures coordination in the use of epistemic procedures across a community, but not rational coordination. It is an essential part of Dogramaci's proposal that epistemic evaluations <u>merely</u> cause, rather than rationally influence, those who are evaluated, to be more (or less) likely to use some epistemic procedure. So, though the members of any given community <u>will</u> (eventually) share a single set of epistemic procedures, no member of the community will have a <u>reason</u> to use the same procedure as his or her fellows. And, by the same token, though all members of the community <u>will</u> trust the testimony of other members, they will have no <u>reason</u> to do so.

We believe that the problems with epistemic communism can be overcome by liberalizing our epistemic economy. Testimony, we would like to suggest, is an excludible good, and its exchange is, in many cases, dictated by market forces. It has value both because the information that it provides can be used in practical or theoretical reasoning ('consumption value'), but also because it can be traded for other goods ('exchange value').[4] Epistemic agents have a practical reason to maintain a supply of testimony with high exchange value, because doing so ensures that they will be able to trade that testimony for other things, most notably other pieces of information, that they may need. We will argue that a properly amended version of epistemic communism does not face the problems that we have identified here.

2. Free Riders

According to epistemic communism, the function of terms of epistemic evaluation is to help promote the coordination of epistemic rule-following across a community. As Dogramaci puts it: 'I pressure you to follow my rules, you pressure me to follow your rules: together we push toward an equilibrium in which we follow shared rules.'[5] There is an advantage to belonging to a community whose members all follow the same set of epistemic procedures, in that it ensures that any testimony that one receives will reflect an epistemic procedure that one endorses. We will argue, however, that epistemic communism does not have the resources to support or promote an equilibrium in which all members of a community share a set of epistemic procedures. Epistemic agents can enjoy the rewards of shared epistemic procedures without paying the cost of policing the epistemic behavior of their fellows. Therefore, if it is assumed that agents will act in their own rational self-interest, epistemic communism will face the same free-rider problems as many other theories of cooperation.

Consider a community with three members: A, B and C. Each talks to the other two, and evaluates their epistemic procedures in the way and with the effect that Dogramaci envisions. Imagine, first, that this is a homogenous community, in which all members share a set of epistemic procedures. If Dogramaci is right about how epistemic evaluations work, all communities should be expected to become homogenous, in this sense, after some number of interactions between their members. Consider this from C's perspective. A and B share a set of epistemic procedures with C, and their interaction provides a check on deviation from those procedures. So C has a reason to trust that A

4

and B will continue to employ epistemic procedures of which she approves, even if she does not evaluate their behavior herself.

Now, there is a cost associated with evaluating someone's epistemic procedures. What someone believes is a product of the information they possess together with the epistemic procedures that they use to process that information. It is often not trivial to determine what combination of evidence and epistemic procedures produced an agent's belief.[6] The time spent identifying and correcting the errors of others constitutes an opportunity cost: it is time that one does not spend pursuing one's own ends.

Since C can trust the testimony offered by A and B without incurring the cost of policing their behavior, it is rational for her to refrain from examining and evaluating them. Of course, C is not special in this respect. For each member of the community, it is rational for that member to count on others to incur the costs of examining and correcting the behavior of other community members. This is the nature of cooperative endeavors such as these, in which there are individual costs and collective benefits; so long as there are benefits that can be had without paying costs, and the magnitude of the benefit is largely independent of whether any given individual participates, it is always rational to be a free rider.[7] So, in homogenous communities, free riding is a dominant strategy, and, in the long run, we can expect homogenous communities to be composed primarily of free riders.[8]

Novel epistemic procedures can be introduced into a community either through the addition of new members, or through innovations produced by current members. Free riding is a problem once novel epistemic procedures are introduced into a community. In such a community, as free-riding goes to fixation, agents will no longer

have a reason to trust the testimony that they receive, because there will be no mechanism to enforce the coordination of epistemic procedures.

In the literature on evolutionary game theory there are a number of solutions to problems of altruistic coordination, such as the problem that we find here. One way that coordination can evolve and remain stable in these circumstances is by means of indirect reciprocity. Indirect reciprocity functions when there is a payoff for aiding others, though that payoff is not provided by those who receive the aid.[9] Indirect reciprocity resolves the collective action problem identified here if there is something valuable that epistemic agents will provide to those who evaluate the epistemic procedures of others.

Call those who regularly evaluate the epistemic procedures used by other members of the community 'epistemic altruists'. On the model that we propose, epistemic agents derive two benefits from establishing a reputation as epistemic altruists. Moreover, these benefits accrue to those who actually are epistemic altruists, because an efficient way to establish a reputation as an epistemic altruist is to actually be one. The first benefit that epistemic altruists secure for themselves is that they increase the exchange value of their testimony. The second is that they increase their access to the testimony of others.

Like all excludible goods, the exchange value of testimony is a function of the supply available, and the demand amongst consumers. The major determinants of the demand for testimony are the degree to which it is useful, the degree to which it is taken to be reliable, and its scarcity. One way to increase the demand for one's testimony, and hence its exchange value, is to increase the degree to which it is taken to be reliable. Since epistemic agents trust the epistemic procedures that they employ, one way to

increase the degree to which one's testimony is taken to be reliable, by those in one's community, is to spread one's favored epistemic procedures in that community. By evaluating other community members, epistemic altruists spread their favored epistemic procedures, or reinforce their use. Hence, evaluating others can drive up demand for the testimony that one produces, as all of those whom one evaluates thereby become potential consumers. So one of the payoffs of being an epistemic altruist is that it increases the exchange value of one's testimony.[10]

The other payoff is related. Imagine a community including four agents, A, B, C and D. Say that A evaluates B and C, and that they come to use alpha, his favored epistemic procedure. If D sees that B and C now use alpha, he can see that there is a benefit to be had by doing so as well.[11] In particular, it will allow him to demand more of A, B and C, in exchange for his testimony, than he could have otherwise.

But it may not be obvious to D how to use alpha, or when it is appropriate to use it. D might be able to figure this out for himself, but he might not. In any case, one efficient way to get the information that he needs is to find someone who already has it, and get it from them. Having corrected B and C, A is <u>conspicuously</u> knowledgeable about alpha, and so is a likely source of information for D. How can D get the information from A? By offering testimony to A, and awaiting correction.

A reputation as an epistemic altruist makes one a likely source of epistemic correction, and hence a likely recipient of testimony that one would not have access to otherwise. If B and C do not enjoy the reputation that A does, D is unlikely to offer them his testimony. This is not to deny, of course, that there are other ways of acquiring the relevant information. For some procedures, perhaps D could figure it out himself. But if

being evaluated by an epistemic altruist is a good way for him to acquire the information that he seeks, then being an epistemic altruist gives A access to testimony that he would not have had otherwise.[12]

Suitable emendation will allow epistemic communism to deal with the problem of free riders. The emendation that is required is an account of the practical reasons that community members have for evaluating the epistemic procedures of their fellows. Epistemic altruists increase the exchange value of their testimony by creating a market for it, and, by creating this market, they provide other members of the community with reason to offer them their testimony. Together, we believe that the advantages that accrue to epistemic altruists are enough to motivate community members to evaluate the epistemic procedures used by their fellows.

3. Epistemic Alienation

Epistemic evaluations can be used to shape *any* epistemic procedure, even the most basic. But one's most basic epistemic procedures cannot be shaped rationally, for a rational influence on one's epistemic procedures would require presenting an argument that one's procedures ought to be modified, but it is by the standards of one's most basic epistemic procedures that any argument would need to be evaluated. So the effect of epistemic evaluations in shaping the procedures that agents use must be non-rational.[13] This can be a problem, because it means that agents who are subject to these non-rational influences will not, thereby, acquire a reason to adopt new procedures, or accept testimony that reflects those procedures. The extent of this problem depends on how disagreements about epistemic practices are resolved. Dogramaci's view is compatible

with both <u>liberal</u> principles for resolving disputes, where both agents modify their procedures following a round of epistemic evaluations, and with <u>conservative</u> principles, where only one agent revises their procedures. We will suggest a conservative principle that helps minimize this problem.

Suppose there are just two agents, A and B, with two conflicting epistemic procedures. Their respective procedures give them a reason to trust their own beliefs, and not the testimony provided by the other agent. Suppose that A moves first, negatively evaluating B. Doing so makes it less likely that B will continue using his epistemic procedure. As he must use <u>some</u> procedure, he becomes more likely to adopt A's. Say that he does.

After one round of interaction, the population arrives at a stable equilibrium, and no further evaluations are necessary. Before the interaction B did not have a reason to trust A's testimony.[14] Does B have a reason after the interaction? No. B <u>will</u> trust A's testimony, but he has been given no <u>reason</u> to do so. A's influence on B is <u>merely</u> causal. He will use A's epistemic procedures, and he will trust A's testimony, but he will have no reason to do so. We will say that B is <u>alienated</u> from the epistemic procedures that he uses, and the testimony that he accepts.

In the case just discussed, only half of the community suffers from epistemic alienation. This was a result of the fact that A moved first, and B had no opportunity to evaluate A. More plausibly, however, A and B will evaluate each other more-or-less simultaneously. Debates about epistemic procedures are not one-sided affairs, and negative epistemic evaluations are often met with resistance.[15]

Say that, before interacting, A used procedure alpha and B used procedure beta. A evaluates B negatively, and B does the same to A. If Dogramaci is right, this means that A will become less likely to continue using alpha, and B will become less likely to continue using beta. It is natural to suppose that disagreements of this kind will be resolved by splitting the difference between the two procedures, or otherwise compromising.[16] Call the new procedure that they settle on 'gamma'. After interacting, A and B both use gamma, but neither A nor B has a <u>reason</u> to use gamma. Similarly, A will trust the testimony that B offers, and B will trust the testimony that A offers, but neither will have a <u>reason</u> to trust this testimony. If epistemic evaluation is not unidirectional, but involves mutual evaluation, we should expect all, or most, members of a community will be alienated from their epistemic procedures and the testimony that they accept. That we should be so alienated is implausible in its own right, but it also shows a weakness internal to epistemic communism, for it indicates that (4) does not follow from (1), (2), and (3).

Solving the problem of epistemic alienation, we believe, requires refining part of Dogramaci's program. Specifically, this will involve reformulating (1) in a more conservative way. He claims that epistemic evaluations tend to discourage or promote the employment of epistemic procedures, depending on whether or not they are evaluated negatively or positively. This effect is a product of the fact that we are 'built to respond to each others' evaluations in accommodating ways.'[17] He goes on: 'Each of us knows from personal experience how we are built to respond to the sting of criticism, or the pleasantness of praise.'[18] If this effect is a product of the bare fact that we seek to avoid criticism and garner praise, one should expect that, if all else is equal, all epistemic

evaluations should have the effects that Dogramaci predicts. We would like to claim, by contrast, that, if all else is equal, the effectiveness of a negative evaluation of some epistemic procedure, *EP,* is directly proportional to the degree to which the community in which the evaluation occurs is united in rejecting *EP.* And, by the same token, if all else is equal, the effectiveness of a positive evaluation of some epistemic procedure*, EP*, is directly proportional to the degree to which the community in which the evaluation occurs is united in approving that procedure.

Say that alpha is widespread in the community in which A and B live. Then when A evaluates B, B becomes less likely to continue using beta, and so more likely to adopt alpha. On the other hand, when B evaluates A, his evaluation is likely to be ineffectual. Eventually, the community will reach a stable equilibrium in which everyone uses alpha.[19] It is obvious, then, that A will not be alienated from his own epistemic procedures (they do not change), nor will he be alienated from the testimony that he receives from B (it reflects epistemic procedures that he has reason to trust). Our modification of epistemic communism has an advantage over the original, because, in most cases, B can <u>also</u> be shown not to be alienated from his procedures, or from the testimony that he receives.

If alpha is widespread in B's community, that fact is evidence that using alpha confers advantages on those who live in the environment in which B finds himself. The strength of this evidence will be proportional to the size of B's community and the length of time that it has survived. So, for example, if there is a small cult centered on its charismatic founder, that some epistemic procedure is widespread in that cult is little reason to believe that it is practically advantageous to use it. But if a particular epistemic

procedure is widespread in a large country with a long history, that fact is rather better evidence that using that procedure is practically advantageous. Moreover, to the extent that epistemic procedures that are practically advantageous to use in a particular environment are likely to be reliable in that environment, one can infer that the community's favored epistemic procedures are likely to be reliable in the environment in which the community lives.[20] So, <u>provided that beta will allow him to infer that using alpha is practically advantageous</u>, B has a reason to adopt alpha. Indeed, he will have two reasons. His inference allows him to infer that beliefs formed using alpha have a high consumption value. And, since alpha is widespread in his community, testimony that reflects the use of alpha will have a high exchange value. So if he is to switch to using alpha he will produce beliefs with a high consumption value, and his testimony will have value in his community. B, therefore, need not be alienated from his new epistemic procedure, nor need he be alienated from whatever testimony that other members of the community provide to him.

We cannot, however, <u>guarantee</u> that B will not suffer from epistemic alienation. Most epistemic procedures that someone might actually use will allow B to infer that using alpha is practically advantageous, from the fact that its use is widespread in a large and well-established community. Say, however, that B's favored epistemic procedure is counter-induction. In that case, he will see that using alpha has been practically advantageous in the past, and will infer that it will not be practically advantageous in the future. Some epistemic alienation may be inevitable. But that there is some alienation is not implausible.[21] What is implausible is that epistemic alienation should be widespread, which is what epistemic communism, if not suitably amended, implies.

4. Summary

As Dogramaci develops it, epistemic communism faces a free rider problem, and it generates epistemic alienation. The solution to these problems is to liberalize our epistemic economy. Testimony has both consumption and exchange value. It is, moreover, an excludible good that is distributed according to market principles. Recognizing these facts allows us to identify incentives that discourage free riding, and to propose modifications to epistemic communism that reduce epistemic alienation.

References

Arkes, H. R. 1991. Costs and benefits of judgment errors: implications for
      debiasing. *Psychological Bulletin*, 110 (2): 486-98.

Brandom, R. 1994. *Making it explicit: reasoning, representing and discursive
      commitment*. Cambridge: Harvard.

Dogramaci, S. 2012. Reverse Engineering Epistemic Evaluations. *Philosophy and
      Phenomenological Research*, 84: 513-20.

Elga, A. 2007. Reflection and disagreement. *Noûs*, 43 (3): 478-502.

Fischhoff, B. 1982. Debiasing. In *Judgment under uncertainty: heuristics and biases*,
      edited by D. Kahneman et al, pp. 422-44. New York: CUP.

Henrich, J. and F. J. Gil-White 2001. The evolution of prestige: freely
      conferred deference as a mechanism for enhancing the benefits of cultural
      transmission. *Evolution & Human Behavior*, 22 (3): 165-96.

Larrick, R. 2004. Debiasing. In *Blackwell handbook of judgment and decision making*,
      edited by D. Koehler and N. Harvey, pp. 316-37. New York: Blackwell.

Nowak, M. A. 2006. Five rules for the evolution of cooperation. *Science*, 314: 1560-3.

[2] Dogramaci 2012: 519-520. Portions of this selection are italicized in the original.

[3] See Dogramaci 2012: 524.

[4] Testimony can be traded for other testimony, but it also plays a role in the broader economy. If Jones is a skilled engineer, prospective employers will find him a desirable employee both because of what he can do, but *also* because of what he can tell other people at the firm. (E.g., that some design developed by another engineer is unlikely to be successful.)

[5] Dogramaci 2012: 522.

[6] See Larick 2004, Arkes 1991 and Fischhoff 1982 for information regarding the challenges involved in shaping the epistemic behavior of others.

[7] It is likely, particularly in large societies, where agents interact freely, that if any given agent refrains from evaluating the epistemic procedures of his or her fellows, that there will be no loss of coordination in that society. The reason for this is that epistemic communities, particularly modern ones, are thoroughly interconnected. If any given agent fails to correct the epistemic behavior of any other, that agent's behavior will likely be corrected by someone else.

[8] How long it will take for free riders to dominate a community depends on contingent facts about its members, concerning, e.g., how rational and well informed they are.

[9] See Nowak 2006.

[10] Seen Henrich and Gil-White 2001 for development of a similar view, according to which an agent benefits from offering instruction to others because this improves the agent's reputation.

[11] It is worth noting that there is some benefit to conformity, even if D takes alpha to be unreliable. Since A, B and C all use alpha, if D's testimony reflects the use of alpha as well, they will find D's testimony more valuable than they would otherwise. By increasing the value that they place on his testimony, D increases the price that they would be willing to pay for it (be it in testimony, or money, or whatever). The benefit to D is in the increased price that he can demand for his information. Of course this benefit might not override the costs associated with changing epistemic procedures. Whether or not it will depends on particular facts about the situation, concerning, for example, what A, B and C have to offer, and whether they could get the information that D can provide from some other source. On the other hand, if D takes alpha to be as reliable as his own procedure, he probably has conclusive reason to switch to alpha, as doing so involves no loss of reliability, but does increase the value of his testimony to others.

[12] We should note that testimony is good to have even if one does not take it to be reliable. If I believe that Smith uses counter-induction (that is, that he infers that *p* will probably be true in the future from the fact that it has been false in the past), and if Smith tells me that it probably will not snow in Boston next winter, I can infer that it probably *will* snow in Boston next winter. Even though Smith is unreliable, his testimony can be

14

mined for useful information. (For more on this theme, see Brandom's account of communication in chapter eight of Brandom 1994.)

[13] Dogramaci is quite explicit that the influence of epistemic evaluations cannot be a rational influence: "An important point requires emphasis here: it's *not* my view that we criticize Smith, Jones and Brown as a way of *reasoning* with them (or with our audience, if the audience is someone else)." (Dogramaci 2012: 521. The emphasis is in the original.)

[14] There may be an "external" reason for B to trust A's testimony, if, say, A's testimony is correct. But it is not a reason had by B, and is not the kind of reason at issue here. We would like to thank an anonymous reviewer for pointing out the fact that the distinction between internal and external reasons may be important here.

[15] Anecdotal evidence for this abounds. *The New York Times* and *The Wall Street Journal* have been known to issue competing editorials on the same day. Democrats and Republicans have been known to hold competing press conferences. Their disputes may be substantive, but they are also often methodological. Disputes among politicians over, for example, whether or not global warming is a result of human behavior are, in part, disputes over whether or not to use scientific procedures when answering questions about the climate. And popular debate – the kind that politicians engage in on the evening news, not the kind that scientists engage in through the journals – rarely goes much beyond issuing evaluations of the kind that play a role in epistemic communism.

[16] There is an obvious connection between this discussion and the literature on peer disagreement. See Elga 2007, for an influential view that requires a compromise between parties to peer disagreements. It is not, however, immediately obvious that Elga's arguments apply directly to the dispute between A and B. Since the influence that they have over each other is *non*-rational, it is not clear that a *rational* means of resolving their dispute is required.

[17] See Dogramaci 2012: 520.

[18] Dogramaci 2012: 520.

[19] In case it is not clear, none of this is to deny that a single individual with a novel epistemic procedure can persuade a community to adopt that practice. Our claim is that the strength of the *non-rational* influence of epistemic evaluations depends on how united the community is in its use of the procedures being evaluated. An individual could still spread a new epistemic procedure by, for example, displaying *evidence* that demonstrates its superiority to those current in the community. Dogramaci's argument concerns only the non-rational effect of epistemic evaluations, and so does our reply. But neither he, nor we, deny that there are other ways to spread the use of an epistemic procedure.

[20] To the extent that advantageous epistemic procedures need not be reliable, B's reasons for adopting alpha will be practical instead of epistemic. But practical reasons are still reasons, and still provide a reason to adopt alpha. Our argument does not, therefore, depend on the claim that advantageous procedures are reliable ones. It would be convenient if advantageous procedures were reliable ones, for then those with novel epistemic procedures have epistemic reasons for adopting the community's procedures. If, inconveniently, it turns out that reliability does not accompany practical advantage, we will be content to say that those with novel epistemic procedures have a merely practical reason to adopt the community's procedures.

[21] Here is a brief argument for this claim: Using counter-induction reduces biological fitness. So individuals who believe in counter-induction yet survive and reproduce probably also use some other epistemic procedures. But counter-induction will sanction the use of only unreliable epistemic procedures. So if it is plausible that there are some individuals who use counter-induction, then it is plausible that there are some individuals who are alienated from some of the procedures that they use. And it is plausible that there are individuals who use counter-induction, at least some of the time. Consider: the gambler's fallacy is in instance of a counter-inductive inference, and there are people who believe in the gambler's fallacy. Therefore it is plausible that some people are alienated from at least some of their epistemic procedures.