

Antisocial Modelling

Georgi Gardiner
University of Tennessee

Forthcoming in *Social Virtue Epistemology*
eds. Mark Alfano, Colin Klein, and Jeroen de Ridder. Routledge

A Reply to Morreau and Olsson's
"Learning from Ranters: The Effect of Information Resistance
on the Epistemic Quality of Social Network Deliberation"

Morreau and Olsson employ a simulation to investigate epistemic effects of false assertions. Ranters are defined as "information-resistant agents that repeatedly broadcast messages conveying their fixed beliefs on the relevant topic", where "information-resistant" means they never update beliefs in light of new evidence. There are two kinds: False ranters only ever convey false claims; true ranters only ever convey true claims. "Openminded" agents, by contrast, update credences in light of evidence.

Morreau and Olsson conclude "Our study suggests that including ranters has little or no negative effect on the epistemic value of social deliberation. Including them can even be epistemically beneficial if the open-minded agents in the network continuously update their trust or distrust in other agents." Even if 80% of agents in the deliberation are false ranters, they claim, "the benefit of deliberation in terms of [their measure of epistemic value] remains about the same ... So, [...] including ranters does not diminish the epistemic value of deliberation for open-minded agents." Call this the "surprising result".

This result relies on openminded agents' ability to track who speaks reliably and to "trust update", that is, to adjust the degree to which others' assertions affect their credences. Openminded agents must even begin to treat false ranters' assertions that p as evidence that not p . The surprising result also relies on their novel interpretation of epistemic value. By "epistemic value" Morreau and Olsson mean "epistemic value for the openminded" (EVO). EVO measures average degree of divergence between credence and truth-value, when restricted to openminded agents. It excludes ranters' credences.

In what follows, I first sketch a concern about the realism of the simulation. I then enumerate some candidates for epistemic value not reflected in EVO and, in some cases, in the simulation more generally.

In the simulation, traits are binary, uniform, and unchanging. Everyone either Bayesian updates flawlessly on all evidence or they never alter any beliefs at all. The latter—ranters—have either comprehensively true or comprehensively false beliefs. All non-social evidence sources emit uniform chance of accuracy, namely 70%. Reality is more complex: A person speaks truly sometimes, and not others, and are more reliable on some topics. People might aptly modulate expressions of confidence, but lapse into overconfidence when discussing politics, or become too diffident around the highly educated. Real-life testimony usually consists of multiple interrelated sentences and so can be partially

accurate. Non-social sources are also heterogenous. Some kinds are more reliable than others. And both inquirers and their sources change—we can become better at interpreting and assessing evidence. This bears on whether the simulation reflects real-life. It is harder to discern, retain, and employ track records when they are complex, evolving, and nuanced, and when we lack dependable or predictable external sources to calibrate against. This variability makes trust updating—an essential component for their result—more onerous and prone to error.

Secondly, the epistemic value of false ranters for openminded agents—the mechanism underwriting the “surprising result”—is that openminded agents treat the false ranters’ assertion that p as evidence that not p . I am doubtful that people do this in real life. It seems contrary to interpersonal interaction and inimical to the institution of testimony. Even considering people I deem most epistemically irresponsible, such as children or anti-vaccine astrologists. I still don’t do it. If they assert some arbitrary claim, I consider it evidence or, at worst, perhaps neutral. I do not consider it evidence *against* p .¹

There are special cases where assertions can be reason to believe the opposite. But they are marginal, require significant background evidence and context, and are about limited domains of assertion. If your acquaintance has contrasting political views, for example, their assertion about the best political candidate can be counterevidence. Tribal, polarised political landscapes can also lead to people taking assertions as counterevidence or, at least, purporting to. Examples include assertions like “Mask-wearing is a safe, effective way to block COVID transmission”. Perhaps we also treat some kinds of aesthetic judgement as counterevidence. We might learn that a person’s musical taste differs so radically from ours, for example, that their liking a musician is evidence that we will not, even if we know nothing else about the musician, such as the genre. But such examples are constrained to limited domains, and do not apply to all assertions.

These two concerns about realism intertwine. We do not treat assertions as counterevidence because real people are not close to uniformly unreliable, like false ranters in the simulation. Perhaps if false ranters lived among us, we would do this. But given realistic heterogeneous track records, I doubt we realistically can “treat what he says as evidence to the contrary” for all or arbitrary-selected assertions of another person. Simulations never aim to perfectly mirror reality, of course, but Morreau and Olsson’s idealisations threaten their conclusion and its applicability to society.

I now turn to costs that Morreau and Olsson do not countenance. I present a miscellany of candidate values and disvalues. Some are measurable in their existing simulations or can be incorporated into new simulations. Others not. Some values accord with veritistic approaches; others not. And, finally, theorists disagree about which candidates are indeed valuable, and whether that value is epistemic. I lack space to investigate these categorisations and questions; I simply enumerate some candidate values.

Morreau and Olsson’s measure of epistemic value, EVO, concerns the average accuracy of openminded agents in the network after a set number of interactions. EVO disregards ranters’ credences. Morreau and Olsson celebrate that EVO remains high even when 80% of group members are false ranters: “[False ranters] develop reputations for speaking falsely, [which makes] their ranting

¹ I am grateful to Catherine Elgin, Liz Camp, and Hilary Kornblith for insightful discussion and for suggesting the political examples.

harmless” because they are not believed. And, since their assertions that p are treated as evidence that not p , “including them can even be epistemically beneficial”.

Morreau and Olsson’s justification for excluding ranters from their measure of epistemic value is twofold. Firstly, ranters don’t learn, so—they claim—their errors are irrelevant to a study of social learning. Morreau and Olsson comment “some people never learn”. But, taken literally, this is false: All people learn. Indeed, since the simulation’s ranters never adjust “beliefs” in light of evidence, and every “belief” is false, I am doubtful the nodes warrant the labels “agent” and “belief”, as opposed to mere information sources. If impervious non-learning justifies excluding them from the measure of epistemic value, it may correspondingly cast doubt on the simulation’s claim to model epistemic interactions amongst different kinds of agents: Ranters aren’t agents.

On the other hand, if ranters are agents, excluding them risks qualifying as creative accountancy: Poorly performing agents are omitted from the record. Discounting them might be elitist, condemning them as deplorables. An epistemic community is better if it improves the epistemic character and conduct of its members, and society should discourage evidence-resistance. In real life, the presence of other delinquents reinforces and sustains epistemic delinquency. A lower incidence of ranting might help rehabilitate ranters. Even setting aside developmental, diachronic considerations, and focusing only on graded alethic accuracy after a number of simulation steps, distribution of accuracy can matter. Non-polarised groups are plausibly better. EVO, as a measure of epistemic value, overlooks these potential disvalues.

Secondly, Morreau and Olsson claim that excluding ranters from the measure of epistemic value better addresses their research question, namely, “how much might unknowingly including [...] incorrigible sources of misinformation hinder your own open-minded search for the truth?” By excluding ranters, EVO better approximates expected epistemic value for “the agents whose epistemic situation might, as far as you know, be your own”.

In response, firstly, openminded people suffer when surrounded by incorrigible ranters. This essay sketches some costs they encounter. So from the perspective of openminded agents, EVO overlooks relevant values. And, secondly, epistemology should also consider an impartial perspective. It should ask not only whether the community is propitious “for me”, but also how the epistemic community is faring overall. If 80% of members have 100% false beliefs, the answer is *badly*.

An epistemic community is better if most beliefs and assertions are true. In the simulations, many assertions are false, disbelieved, or treated as counterevidence. A prevalence of true assertions has instrumental epistemic value. Discerning and recalling who is unreliable demands cognitive resources, which has opportunity costs. Trust updating requires dynamic accountancy. One must track whose previous assertions conflicted with one’s antecedent beliefs and weight their future testimony accordingly. The effort is better invested elsewhere. And trust updating is fallible. Learning is easier if we can generally believe people.

Indeed, absent background knowledge about the bifurcated epistemic community—that is, ranters and others—encountering false ranters in the simulation should decrease trust in testimony simpliciter. Testimony appears—and is—unreliable.

A prevalence of true assertions may also exhibit non-instrumental epistemic value. There is plausibly cognitive value in attention being directed towards the right things. Perhaps proper attentional patterns

are constitutive of flourishing, for example. This is why Aristotle posits that God's sole activity is self-contemplation. Untangling a morass of false assertion, by contrast, is a lousy activity to absorb attention. Well-functioning trust relationships may also have non-instrumental value. Systems are better when they function properly, rather than deviantly. Treating assertions that p as evidence against p is dysfunctional.

Some disvalues are emergent. Dyads of distrust are bad; a prevalence is corrosive. Social institutions would dissolve; many would never have arisen. This likely includes the institution of testimony itself. Discussion allows people to develop skills, cultivate virtues, expand human understanding, perceive gaps in human knowledge, and develop appropriate humility. Learning together is more than adjusting confidence in isolated propositions. We help one another understand and interpret the world. Together we forge new conceptions and formulate better questions. We can be role models and inspire each other epistemically. By learning together, communities bond. These values are threatened by widespread ranting, evidence-resistance, and a prevalence of distrust.

Morreau and Olsson conclude,

Excluding information resistant people from debates requires recognizing them as such, which realistically must consume resources. Censorship, de-platforming and other ways of excluding them furthermore violate social and democratic values of inclusiveness and free speech. Our study suggests that, in some contexts anyway, the epistemic benefits of excluding false ranters are not worth these costs. [...] It might be better simply to let everyone have their say.

But tracking false ranters, ignoring assertions, and treating testimony as counterevidence also have significant epistemic, social, moral, and opportunity costs. To apply simulations and formal measures of epistemic value to real life, these costs cannot be ignored.²

Acknowledgements

I am grateful to Catherine Elgin, Liz Camp, Hilary Kornblith, and Cat Saint-Croix for helpful comments, and to Jon Garthoff for discussion and feedback on an earlier draft. This research was supported by an ACLS Fellowship from the American Council of Learned Societies.

² Morreau and Olsson write, “there are related studies suggesting [...] that restricting social information, or *equivalently*, [individuals] being resistant to such information can sometimes be epistemically beneficial” (emphasis added). But restricting information (that is, censorship and similar) and an individual's not updating on that information are not equivalent. They constitute different socio-epistemic arrangements, require different epistemic institutions and cognitive resources, and generate different counterfactual epistemic conditions. If a simulation and measure of epistemic value cannot distinguish them, this redounds poorly on the model.