

New frontiers in epistemic evaluation: Lackey on the epistemology of groups

Groups are intriguing targets of epistemic evaluation, as Jennifer Lackey's new book on group epistemology makes abundantly clear: familiar questions about knowledge and justification take on a stimulating new level of complexity when we consider their application to social entities such as teams, schools, and corporations. Groups are also important targets for these assessments, given their power. The power of groups gives the book a sense of urgency: Lackey deftly underscores our need to hold groups accountable for what they have done, which in her view will demand evaluations of what they know and believe.

To appreciate both the importance and the difficulty of her project, one might start by situating the epistemology of groups in its larger natural context. There is a deep relationship between the human capacity for epistemic evaluation and the human capacity to form groups. Across the animal kingdom, many species have specialized forms of intelligence for dealing with other agents, in their roles as conspecifics, predators, and prey. In humans, this social intelligence has a radically distinct character. On top of our basic animal sensitivity to the sightlines and goals of others, we have a deeper appreciation of what is driving them. In particular, we seem to be the only species capable of attributing beliefs: while many animals can grasp what others see, and what is hidden from their view, humans can track and evaluate more subtle mental states, including states of mind whose contents may deviate from our shared reality (Rakoczy, 2022; Tomasello, 2018). This powerful sense of other minds has given our species a deeper capacity to cooperate in groups. Social insects cooperate seamlessly, but only in a genetically prescribed manner, adhering to fixed social roles and sharing information restricted in advance to a small set of topics such as the location of nectar. We cooperate with each other in novel and flexible ways, sharing information about a boundless variety of topics, correcting each other's misconceptions, and setting new goals reachable only by combining our forces. Michael Tomasello has influentially argued that this human ultra-cooperativeness is what accounts for the astonishing distinctiveness of our species: our capacity to work together in novel groupings has essentially created a new level of evolution, in which socially generated structures enable the creation and transmission of abilities and knowledge beyond the reach of any individual, an effect which ratchets up from one generation to the next (Tomasello, 2019).

As groups ranging from schools to corporations have grown in scope and power, it is natural for us as individuals to apply our social intelligence to these novel social entities themselves, for example to wonder what Google knows about us, or what General Motors really believes about climate change. This turn, however, raises questions about the legitimacy of applying our natural mindreading instincts, evolved for face-to-face interactions with individuals in small scale societies, to these new larger-scale social formations. In this context, Jennifer Lackey's incisive analysis of the epistemic states of groups has a lot to offer. Anyone who is worried that it is at best metaphorical to see groups as having beliefs can gain a new appreciation of group doxastic powers from Lackey's rich set of examples. She develops an intriguing model of the conditions under which groups properly count as believing, being justified, and knowing, a model guided above all by the thought that we need to explain how groups can be held responsible for what they do. She supports this model largely with the method of cases, identifying situations in which it seems entirely natural to describe a jury, corporation or other group as holding some belief, or guilty of lying, often situations in which

there might be an appropriately parallel (and less controversial) case of an individual agent who has a belief or tells a lie. If our core understanding of epistemic mental states (and their abuses) comes from interactions with individuals, as seems entirely plausible, then this method of drawing out parallels with the simpler individual case makes abundant sense, assuming the individual case really is uncontroversial. In what follows, I raise a few questions about particular points at which I think perhaps the lesson for the group case from the level of the individual is not as clear as Lackey might want. Then, at the end, I will raise a slightly larger question about some recent developments in group belief, developments of a type that may present a deeper challenge to our natural capacities for epistemic evaluation, honed as these capacities are on individuals who are roughly our peers.

My first small questions are about Lackey's arguments against inflationary views of group belief, in which group belief is not a straight function of the individual beliefs of the members, but for example the product of something like a decision on their part to accept something for group purposes. In her view, these views err in allowing groups to count as believing propositions jointly believed or accepted by the group in cases where members of the group believe or accept this proposition on frankly incompatible bases. In Lackey's view, it is a condition on the attribution of a belief that it "cannot be base fragile, where this means that the bases of the individually-held beliefs are wildly conflicting" (p.46). She reasons as follows: "even if every member of a group believes that *p*, they might do so for wildly conflicting reasons, which renders the resulting state unfit for epistemic evaluation and for future deliberation in relation to group action. This, then, prevents the state from being a group belief."

Lackey is certainly right that something unsavory has happened in a case of base fragility, but I wonder whether we must take all such cases right off the table of epistemic evaluation. We might, for example, want to say that some anti-vaccination group believes that a major pharmaceutical company's vaccine does not prevent serious illness, where half the members of the group say so on the basis of a belief that the company has faked data of efficacy against a real and serious disease for what is in effect overpriced saline solution, and the other half say that the disease widely taken to be a serious illness is not really harmful, impressions to the contrary being driven by corrupt media (this faction might even grant that the company's product works against its stated target, while firmly denying that this target ailment is a serious illness as opposed to a minor inconvenience). Notwithstanding these clashing grounds, there's a case to be made for the claim that the group as a whole strongly believes that the vaccine does not prevent serious illness: that shared commitment can robustly predict and explain significant behavior on the part of the group, including many of the group's natural inferences to further beliefs, such as the belief that government spending on the vaccine is a waste of money. Crucially, this problematically supported belief could be a factor in the group's cohesion, in the sense that an individual's denying it could be expected to lead to their ostracism or ejection from the group. Where agreement on a point is what unites a group, it is hard to see that agreement as nothing more than a collection of individual commitments. And we might very much want to hold this group responsible for the actions they take on the basis of their shared commitment, however strange its basing.

Meanwhile, rather than disqualifying the group's agreement on this point from epistemic evaluation, perhaps the incoherence of the basis within the group could simply be a reason to evaluate this group state negatively, as an irrational belief. There is indeed a sense on which this belief is unfit for future deliberation: at least some group deliberations invoking the problematically supported proposition could expose the rift between the factions and leave the group divided over what to do or infer next. But one would not want an account of belief which disqualifies from that status any

state that stands at risk of being undermined by future deliberations. At the individual level, many of our irrational beliefs have this status, and attributing these beliefs to us can still explain various decisions we make, sometimes including our decisions to avoid deliberating in those sensitive areas. Lackey points out that new evidence could have a strange impact on base-fragile commitments: “when viewed in light of one set of bases, the evidence counts against the belief, but when viewed in light of another set, it counts in favour of it.” (p.46) In her view, “This deep lack of unity reveals that the state that is purported to be a single one belonging to a group is in fact a collection of individual beliefs.” However, some version of this phenomenon of ambivalent new evidence can arise already at the level of an irrational individual who draws on clashing compartmentalized beliefs to support some further belief. As individuals, having bases for our beliefs that are not substantively incoherent is a rational ideal, not a necessary consequence of what it is to be an agent who can be held responsible. I imagine that at some level, internal incoherence might draw both agency and responsibility into question, but it would be good to hear more about what sets that level, especially if we want to consider the likes of QAnon to count as groups, indeed as groups responsible for the harms generated by their beliefs, notwithstanding the peculiarity and diversity of their basing.

Lackey’s second major argument against inflationary views of group belief is that such views make it too easy for groups to evade responsibility for what they are doing. If belief is a matter of voluntary coordinated acceptance on the part of a group’s members, a group can strategically accept claims its members know to be false in order to enable the group to assert falsehoods with sincerity. If board members at Philip Morris who are individually aware of the harms of smoking nevertheless decide that the corporation will officially accept that smoking is harmless, then Philip Morris will not be lying when it disseminates that view to the public, on the assumption that corporate beliefs are set by deliberate board agreement. The sense that Philip Morris would still be culpable of deception here is a strike against inflationary views of group belief, Lackey suggests. (Conceivably, inflationists might push back with some thoughts about epistemic conditions on proper acceptance, but for present purposes we can set that aside.)

Lackey’s path out of this problem is to devise an account of group belief that tethers group epistemic states to the beliefs of their members:

Group Agent Account: A group, *G*, believes that *p*, if and only if (1) there is a significant percentage of *G*’s operative members who believe that *p*, and (2) are such that adding together the bases of their beliefs that *p* yields a belief set that is not substantively incoherent. (pp. 48-9)

The first clause does not demand unanimity: in her example, an academic department can believe something despite the presence of a single skeptical holdout in their ranks. But the demand for belief on the part of just a healthy percentage of operative members suggests another way that Philip Morris might aim for sincerity: they could decide to expand the size of their board to pack it with true believers. The original board members who knew about the harms of smoking could perhaps outnumber themselves with cheerfully unaware Marlboro men to the point where they would have the status of the isolated skeptic whose contrary view fails to count against the group belief; if the ratio still needs work, surely some old-timers could be offered bonuses for stepping down.

There’s a case to be made that Lackey’s view gets this reformed board right: now Philip Morris really does believe that smoking is harmless. The company might not be lying at this point, but its responsibility is by no means diminished, and arguably this responsibility has a peculiarly epistemic

character. The company has succeeded in some kind of self-deception, turning itself into a true believer of what it once knew was a lie, deliberately structuring itself to press a corresponding deception upon the public. It's not clear to me how Lackey's Group Agent Account handles the epistemic evaluation of this case, or any case in which groups strategically select their future members on the basis of their pre-existing individual doxastic states. If the powers of a corporation can be applied to worsening its epistemic predicament over time, then corporate responsibility is not simply a function of a corporation's current epistemic and motivational states. But now we might wonder about Lackey's overall strategy: if roughly equivalent responsibility persists through a complete change of belief, then intuitions about responsibility might not be the best guide to the nature of belief.

Focusing a different hypothetical variant of Lackey's tobacco company case, in which the board simply sets up procedures to avoid ever collectively pronouncing on the harms of tobacco, Richard Pettigrew offers an analysis of the situation which could apply here as well:

The corporation is liable because the epistemic structures and procedures it's put in place are inadequate to providing them with beliefs that will ensure the safety of their consumers. They are guilty of what Hagemann & Grinstein (1997) call wilful blindness in their critical examination of *United States v. Bank of New England, N.A.*, one of the most widely discussed cases in which the presence or absence of corporate belief was important for the charge of corporate liability. (Pettigrew, MS, p.16)

Legal scholars reflecting on the case cited here have argued that ignorance does not even need to be wilful to ground moral responsibility: in some situations, reckless ignorance can also serve this role (Sarch, 2016). Where groups have group-level motives to be recklessly ignorant, for example because their corporate dividends would be threatened by various types of knowledge, they can select members who simply lack certain potentially problematic knowledge or belief, but as Pettigrew points out, this isn't a way of evading liability or responsibility.

What now might be more worrisome for Lackey is the possibility that her view might admit that Philip Morris *justifiably* believes that smoking is harmless, at least on some possible understandings of evidence and justification. Her account of group justification consists in the following two conditions:

- (1) A significant percentage of the operative members of G (a) justifiedly believe that p , and (b) are such that adding together the bases of their justified beliefs that p yields a belief set that is coherent.
- (2) Full disclosure of the evidence relevant to the proposition that p , accompanied by rational deliberation about that evidence among the members of G in accordance with their individual and group epistemic normative requirements, would not result in further evidence that, when added to the bases of G's members' beliefs that p , yields a total belief set that fails to make sufficiently probable that p .

Let's assume that our Marlboro men are selected for their naïveté: they have no prior evidence of the harms of smoking. Or perhaps they are given extensive evidence of the harmlessness of smoking, through a strategy like the Tobacco Industry Research Committee's, which funded a vast number of small independent studies of the impact of smoking on health in the 1950s, and heavily publicized just those statistically outlying studies which happened to turn up good news for them

(Oreskes & Conway, 2011). Lackey observes, reasonably enough, that “because groups can often exercise control over what is jointly accepted, they can manipulate what evidence is, and is not, available to them as a group” (p.100). Her antidote to this problem is to situate justification at the level of the individual: “if the justification of group beliefs is necessarily a matter of the justification of the beliefs of individual members, and the evidence that is available to individual subjects is not a matter of choice, then there is no worry that epistemic justification for group beliefs can be achieved through the illegitimate manipulation of evidence.” The loophole here is that the evidence that is available to individual subjects might indeed not be a matter of their individual choice, but might be chosen or trained by operative members of the group in ways that will amount to epistemically problematic manipulation. For example, it could set their priors to levels at which they can be expected to be skeptical of the truth, when it does cross their radar. And it is not hard to see how extensive deliberation between these yes-men on the board (even with a small, more knowledgeable minority) could indeed fail to result in further evidence that would fail to make it sufficiently subjectively probable to a significant percentage of them that smoking was harmless.

There are doubtless various ways that the “full disclosure” clause might be construed to get over this problem, and of course many ways of understanding “justification” that would eliminate the possibility of a justified individual Marlboro man in the first place, but I am unsure which way of solving the problem most naturally fits with Lackey’s larger view.

Now to more recent developments on the group landscape. I think it is fair to say that *The Epistemology of Groups* is somewhat old-school in the groups it highlights: academic departments, a traditional tobacco company, courtroom juries, and museum guards. These are significant social entities, and Lackey’s theory does much to illuminate their inner workings. For these groups in particular, it is easy to see power vested in a clear set of operative members, and to trace the epistemic status of the group ultimately to patterns in the thinking of those individual humans. But one might also wonder about newer social configurations – mainly corporations, but also some groups founded on social media platforms – where the operative member analysis is harder to work out.

What I have in mind is the following: I think Google knows where I had dinner on the first night of the last conference I attended. I used Google Chrome to search for a restaurant, I used Google Maps to navigate there, after clicking through and reading some of the reviews, and then, after staying put at that location for roughly 90 minutes, I used Google maps again to navigate back to the hotel. Google didn’t see me actually putting a fork in my mouth, but it didn’t have to, Google knew just what I was up to. (Arguably, Google also already knows where I’m going to be having dinner tomorrow night, and perhaps even who I’m eating with, given my web browsing and the state of my Google Calendar, and potentially also their browsing and the states of their Calendars as well—to say nothing of how much they know about the lives of those who use Gmail). I don’t see that a significant percentage of the operative members of Google have any belief about where I ate. I didn’t interact with any humans on their team, and I’m quite sure my data weren’t brought to anyone’s attention; in fact, I believe it would be a violation of Google’s own privacy standards if anyone were to have taken such a God’s-eye peek at me as an individual. But Google took note of my whereabouts that night nonetheless, and its interactions with me have shifted as a result.

By contrast, if I paid cash for coffee that morning at the hotel Starbucks, I actually met one of the operative members of that establishment, but I can’t say that Starbucks knows I had coffee there; I

expect my barista has already forgotten me, and nothing will change at that outlet, or at the chain, as the result of my purchase.

In my view, the sophistication of the algorithms driving personalized data collection are part of what makes it appropriate to attribute mental states to Google as a corporation. Google is an exceptionally intelligent agent, exercising social cognition of its own in its dealings with us mere humans, constantly monitoring our uptake of its recommendations and tweaking its search results to steer us better. Google has ways of telling what we want, and ways of learning what we are likely to do, as a result of what we have seen. The algorithms it employs, and the vast amount of data it works with, give it a new way of making judgments about what is happening in the world, judgments that shape its subsequent intelligent behavior (and ours) in ways that the humans in the corporation themselves may not fully understand (to say nothing of end users like me). At least some of the epistemic powers of this corporation are something of a black box, and it is not easy for me to see how to allocate them to operative members of the group. One option would be to take a larger view of “operative members”, a view in which algorithmic systems like Google Maps could count as epistemically individual members, whose integration with Google Search, Calendar, and human board members at Google or Alphabet could be a topic of further exploration.

At the same time, one might wonder whether we as humans have managed to create a new kind of agent here, one whose workings we don't fully understand, even as we spontaneously trust what it tells us, where this trust is secured in part by how well it knows us. If so, it seems to me an open question how well our intuitive ways of assessing and evaluating individual doxastic states will apply to agents of this new kind. If our intuitive understanding of mental states has emerged under the evolutionary constraints relevant to understanding fellow individual members of our species, who think roughly like we do, then are we naturally well-equipped to reason about what an entity like Google knows, or do we have to start work on devising some new way of understanding such creatures? I am unsure.

Finally, one more arguably novel type of group is the social media group, whose members are drawn by some shared topic of interest, and then perhaps funnelled into holding increasingly extreme beliefs, perhaps simply by the dynamics of an algorithm set to maximize platform engagement (Bryant, 2020). Here again we see some unusual epistemic power on the part of the algorithm itself, in terms of increasing restrictions on the topical evidence that individual members are exposed to; we can also worry that the epistemic dynamics here function in ways that no individual human really understands. We see group dynamics that take on arguably novel shapes, as individuals are clustered with increasingly like-minded individuals, and perhaps start to believe in order to belong, in ways that might stretch the mind-to-world direction of fit that Lackey takes to be definitive of belief. The power of the algorithm to attract and repel group members seems to give that algorithm an organizational status in the group that again lifts it above the individual epistemic power of any operative member, if operative members always have to be individual human beings. Again, we can worry that our natural capacities for mental state attribution could be challenged by these new configurations of group mentality.

Of course, there's a question about whether the apparent doxastic formations in these groups really should count as beliefs (Mercier, 2020). And doubtless there are questions about whether these groups really are as novel in their epistemic makeup as they might seem: perhaps there is some good way to apply the operative member analysis here, despite the superficial novelty.

There's a reason for my final focus on corporations like Google and social media groups. The leading idea of *The Epistemology of Groups* is that we need to be able to hold group agents responsible, and it seems appropriate to close by expressing warm enthusiasm for that idea, and underscoring its urgency with respect to these strange new group agents, given their rising power in the world. I am quite sure that a large part of holding them responsible will demand making sense of their epistemic states, and I have no doubt that Jennifer Lackey's innovative work on group epistemology can help us to find the right path forward.

References:

- Bryant, L. V. (2020). The YouTube algorithm and the Alt-Right filter bubble. *Open Information Science*, 4(1), 85-90.
- Mercier, H. (2020). *Not born yesterday*: Princeton University Press.
- Oreskes, N., & Conway, E. M. (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*: Bloomsbury Publishing USA.
- Pettigrew, R. (MS). Believing is said of groups in many ways (and so it should be said of them in none).
- Rakoczy, H. (2022). Foundations of theory of mind and its development in early childhood. *Nature Reviews Psychology*, 1-13.
- Sarch, A. (2016). Beyond Willful Ignorance. *University of Colorado Law Review*, 88, 97-178.
- Tomasello, M. (2018). How children come to understand false beliefs: A shared intentionality account. *Proceedings of the National Academy of Sciences*, 115(34), 8491-8498.
- Tomasello, M. (2019). *Becoming human: A theory of ontogeny*: Belknap Press.