




ARTICLE

Defending Joint Acceptance Accounts of Justification

Lukas Schwengerer 

University of Duisburg-Essen, Essen, Germany
Email: Lukas.Schwengerer@uni-due.de

(Received 27 May 2020; revised 7 September 2020; accepted 16 November 2020)

Abstract

Jennifer Lackey (2016) challenged group acceptance accounts of justification by arguing that these accounts make the possession of evidence arbitrary and hence lead to illegitimate manipulation of the group's evidence. She proposes that the only way out is to rely on the epistemic propriety of the individual group members, which leads to a dilemma for group acceptance views: either they are wrong about justification, or they cease to rely only on group acceptances. I argue that there is a third option based on general expectations of epistemic propriety that restricts the group's maximal justification. A group cannot be more justified than any individual in the group's position could be expected to be. I motivate this solution by a discussion of normative defeat and epistemic expectations as proposed by Goldberg (2018).

Keywords: Collective justification; social groups; collective intentionality; collective belief; normative defeat; Jennifer Lackey; Sanford Goldberg

Introduction

Jennifer Lackey challenged group acceptance accounts of justification by arguing that these accounts make the possession of evidence arbitrary and hence lead to illegitimate manipulation of the group's evidence. She proposes that the only way out is to rely on the epistemic propriety of the individual group members, which leads to a dilemma for group acceptance views: either they are wrong about justification, or they cease to rely only on group acceptances. I argue that there is a third option based on general expectations of epistemic propriety that restricts the group's maximal justification. A group cannot be more justified than an individual in the group's position could be expected to be. I motivate this solution by a discussion of normative defeat and epistemic expectations inspired by Goldberg.

My plan is the following: I start with a quick discussion of joint acceptance accounts of group justification in section 1. I then explain Lackey's argument against these views and discuss her suggested response in section 2. In section 3 I draw a parallel to the individual case and normative defeat which motivates my response to Lackey's argument. In sections 4 and 5 I develop my proposal in more detail and relate it to general expectations as part of epistemic evaluation. Finally, in section 6, I discuss potential

© The Author(s), 2021. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

objections to my proposal and show how epistemic expectations towards individual and groups relate.

1. Group justification as joint acceptance of reasons

Accounts of group justification can be roughly grouped into inflational, non-summative approaches that take group justification to be over and above or otherwise distinct from the member's justification, and deflationary, summative views that hold that group justification is merely an aggregation of the member's justification. Lackey recently argued that both of these approaches lead to problems, and proposed a middle ground position. For the purpose of this paper I am only interested in a particular type of inflational, non-summative position that takes group justification to be based on reasons the group accepts collectively. 'Collectively' here is meant to indicate an acceptance as a group that cannot be reduced to attitudes of individual group members. Call this the *joint acceptance account* (JAA) of group justification. Two versions of JAA can be found in Schmitt (1994) and Hakli (2011). I use the slightly modified formulations from Lackey (2016).

JAA-S: A group *G* justifiably believes that *p* if and only if *G* has good reason to believe that *p*, and believes that *p* for this reason, where *G* has a reason *r* to believe that *p* if and only if all members of *G* would properly express openly a willingness to accept *r* jointly as the group's reason to believe that *p*. (Lackey 2016: 346)

JAA-H: A group *G* justifiably believes that *p* collectively "if and only if the group can successfully defend *p* against reasonable challenges by providing reasons or evidence that are collectively acceptable to the group and that support *p* according to the epistemic principles collectively accepted in the epistemic community of the group. The epistemic community determines what counts as a successful defence and as a reasonable challenge." (Lackey 2016: 348)

Both examples of JAA function in a similar way. They propose that group justification can be spelled out in virtue of reasons for a belief that *p* that the group has.¹ Reasons can be epistemically evaluated insofar as they support the belief in question. Reasons to believe that *p* are good reasons if they provide sufficient evidential support, either by being accepted as responses to challenges to the belief by an epistemic community (Hakli), or in virtue of reliably leading to truth (Schmitt). The precise method of evaluation does not matter for now, but will come up later again. To understand Lackey's attack on JAA another feature is more important: the formation of these reasons. Both good and bad reasons are formed in the same way: the group collectively accepts a reason *r*. To do so the group needs nothing more than each of the relevant group members being willing to accept *r* as a reason for the group and – on some versions – the members have to express this willingness openly. In the simplest version that includes such an expression the group members indicate their willingness explicitly in speech. Suppose we have a group of three members: Reid, Melissa, and Kathleen. Reid says that he is willing to accept that *r* as a reason, Melissa answers that she is willing to do the same, before Kathleen states "If you both want to accept *r* then I also will." The three agree to accept *r* as a reason for the group.²

¹Remaining neutral on the nature of reasons. Joint acceptance accounts are compatible with (but not limited to) a mentalist conception of reasons, or a propositional conception of reasons.

²Gilbert holds that this is a kind of conditional commitment. Only when every member accepts *r* the commitment is established. If only Reid and Melissa would have agreed to take *r* as a group reason nothing

The condition that every member has to be a part of the joint acceptance procedure is surely too strong. In some cases particular members are not involved in forming group states, or making group decisions. Reid might always defer to the other members and just execute the decisions made by the rest. Following Tuomela (2013: 86–7) I call those relevant for forming group states *operative members*.³ We can now define the key feature of the formation of reasons under JAA:

A group has reason *r* iff all operative members are willing to take a reason *r* as a group reason (and expresses this willingness openly).

This expression of willingness to jointly accept *r* can be either explicit, or implicit in the member's behaviour. Importantly, each member is willing to let *r* stand as the group's reason, even if they do not accept *r* on an individual level. Melissa can be part of the group's joint acceptance of *r*, even though she personally does not accept *r* as a reason whatsoever. This willingness is a hallmark of group states under non-summativist accounts (Gilbert 1989: 289) and is used to explain cases in which the group state supposedly diverges from the member's states. One such case is provided by Kallestrup (2016), who modelled the case after a discussion in Schmitt (1994).⁴

CRIMINAL COURT: A defendant is on trial for the crime of careless driving. The prosecution adduces evidence from the police report, as well as eyewitnesses testifying in court that the defendant was indeed driving the van that hit the victim. But the jury finds the evidence not beyond a reasonable doubt, and hence insufficient to validate criminal conviction. All the members of the jury have hearsay evidence from a reliable source that the defendant caused the accident. But the judge instructs the jurors to ignore this evidence as it fails to meet the conditions for being ruled admissible in a criminal court. Consequently, the jury, when functioning in its social office, justifiably believes the defendant is innocent, yet none of the individual jurors justifiably believes this proposition.⁵

The jury is a *chartered* group, found with a particular aim and in order to perform a particular action. This aim is the group's *office*. The members have to intend to fulfil the aim, but it need not be made explicit. In this case the group is formed with the aim to come to a verdict regarding the defendant's guilt based on evidence admissible in court. This aim affects what the group members accept as evidence that can justify beliefs on the group level. While each individual member of the jury might believe that the defendant is guilty, they are fully aware that their beliefs are based on evidence that ought not to be taken into consideration for the jury's verdict. The group will not accept

would have happened (cf. Gilbert 1990). Joint acceptance needs all three of them. I want to remain neutral here.

³A reviewer notices that the idea of operative members combined with accounts that do or do not require open expression of the willingness to accept *r* as a reason gives us some specific scenarios that are worth a discussion for which I have unfortunately no room. For instance, if only operative members accept *r* as a reason, but they need no open expression we end up with groups that accept *r* as a reason for *p*, but which includes non-operative group members who do not know that the group has reason *r*.

⁴Lackey (2016) also discusses a version of this case.

⁵There is an alternative reading holding that social roles play a more important role here than the group qua group does. Divergence cases could also be discussed under this aspect. And perhaps the same lesson I suggest for group justification can be generalised to justification a person has as occupying a specific social role. My topic is group justification, so I am not following this thread further. Thanks to Miguel Hoeltje for emphasising this point.

the hearsay evidence as a group reason to believe that the defendant is guilty. Hence, the group will find the defendant not guilty. Lackey (2016: 354–7) raises some objections against cases like CRIMINAL COURT, but for now I want to let the case stand as motivation for JAA. The group members decide on a reason for the group by jointly accepting or not accepting propositions. In this case, the members are not willing to take the hearsay evidence as a reason. The key feature of this reason-formation process is that the group's reasons are constrained merely by the operative member's willingness to accept some proposition *r* as a reason. Groups can form reasons at will, according to JAA. This is the feature that Lackey picks out to generate potential problem cases.

Before presenting Lackey's arguments it is helpful to motivate the aim at hand: why defend JAA? The primary reason is the possibility of divergence cases, such as CRIMINAL COURT. If one accepts any divergence case, then one needs a way to explain group justification that is compatible with no group member being justified. Fully aware of the issue Lackey argues against some divergence cases, including CRIMINAL COURT. However, other divergence cases with different structures are unaddressed: Hakli (2011: 120–1) introduces his view by pointing to a weather forecast group in which the group accepts one collective forecast that no group member holds to be true. Bird (2010: 34–5) suggests that scientific groups can produce knowledge without the members having justified beliefs. Tollefsen (2009: 12) points to a case of the U.S. National Academies of Science committing to a position as a group, even though the group members disagree on the issue in question. One might read this also as a case of justified group belief. Lackey's explanation of CRIMINAL COURT points to a difference between evidence in law and epistemic justification. It is not clear that a similar answer can be given to these other cases.

Of course, all of these cases can be contested. However, as long as one of them is a live candidate for a genuine divergence case the JAA is still a promising candidate for an account of group justification worth defending. My aim is to find a version of JAA that is compatible with divergence cases, but nevertheless not susceptible to Lackey's arguments against JAA. Whether Lackey's view or my version of JAA is ultimately better depends on a final judgment of divergence cases, which I do not provide. I merely want to find a version of JAA that is a live option for a successful theory of group justification. I am not arguing that it is the only account that can do that, but rather that we should not rule it out yet.

2. Lackey against JAA

Lackey raises two different cases illustrating the consequences of forming reasons at will, which she takes to point to "a decisive objection" (Lackey 2016: 350) against JAA:

IGNORING EVIDENCE: Philip Morris is one of the largest tobacco companies in the world, and each of its operative members is individually aware of the massive amounts of scientific evidence revealing not only the addictiveness of smoking but also the links it has with lung cancer and heart disease. Moreover, each individual member believes that the dangers of smoking give the company a reason to believe that warning labels should be placed on cigarette boxes. However, because of what is at stake financially and legally, none of these members would properly express a willingness to accept that the dangers of smoking give Philip Morris a reason to believe that it should put warning labels on cigarette boxes.

FABRICATING EVIDENCE: Philip Morris is one of the largest tobacco companies in the world, and each of its operative members is individually aware of the massive amounts of scientific evidence revealing not only the addictiveness of

smoking but also the links it has with lung cancer and heart disease. Entirely because of what is at stake financially and legally, however, each of these members decides to jointly accept that all of the scientists working on the relationship between smoking and health problems are liars. Given this, they also jointly accept that the duplicity of the scientists gives Philip Morris a reason to believe that the results of the studies showing a connection between smoking and lung cancer and heart disease are unreliable. (Lackey 2016: 351)

In IGNORING EVIDENCE we are supposed to come away with the intuition that Philip Morris clearly has a reason to believe that it should be putting warning labels on the cigarette boxes. Their deliberate ignorance should not result in them lacking such a reason. But JAA claims just that! I agree with Lackey's assessment that this is an unacceptable result.

FABRICATING EVIDENCE goes one step further. Philip Morris does not merely ignore evidence deliberately, but generates reasons for purely pragmatic goals. Even if the members have individually epistemic reasons to believe that the scientists are trustworthy, as a group they are willing to accept that these scientists are liars and the studies not to be reliable. The intuition pumped is that this joint acceptance is not epistemically justifying anything. Clearly, financial incentives should not be able to create epistemic justification for a group out of thin air. Lackey's intuition seems obviously true to me. And it is a problem for JAA, because JAA is compatible with exactly that: forming epistemic reasons at will. JAA does not have any constraints on the process of creating joint acceptance besides all operative members being willing to accept *r* as a reason. For JAA it does not matter why they are willing to accept *r* as a reason. More bad news for JAA!

At first sight, the problem seems to be unavoidable for JAA. These accounts are created with the aim of accommodating divergence cases, like CRIMINAL COURT. They do so by claiming that there is no connection between reasons that individual members have, and reasons the group has. But as soon as one has severed this connection the group's grip on the world is lost. Groups can accept whatever the members feel like accepting. They are not accountable to compatibility with individually held states. And as long as one does not want to propose a group mind that can access the world independently of its members it seems difficult to imagine anything that makes sure that the group's reasons are properly connected to the world. They are arbitrary, and arbitrariness is not compatible with epistemic propriety. Epistemic propriety includes aiming at truth, after all. We end up with Lackey's *Illegitimate Manipulation of Evidence Problem (IMEP)*:

IMEP: If the justification of group beliefs can be achieved through wholly voluntary means, then the evidence available to the group can be illegitimately manipulated, thereby severing the connection between group epistemic justification and truth-conduciveness. (Lackey 2016: 353)

Is there a way out? Lackey claims the only option available is to leave JAA behind. To do so one can introduce a constraint on the individual member's willingness to accept a reason *r* as a group reason. As is, JAA treats every instance of a member's voluntary acceptance equally. But it need not do so. Lackey suggests that some acceptances are epistemically proper, whereas others are not. And this propriety depends on the epistemic situation of the individual member. Plausibly, if a member has a reason *r*, then it is proper for the member to be willing to accept *r* as a reason for the group. Moreover, if the group members disclose their evidences to one another, the group

can base their joint acceptance on the total evidence of the individual group members. Truth-conduciveness of group justification is salvaged in virtue of the truth-conduciveness of the individual members. But there is a price to pay. If the group justification is based on the justification of the group members, divergence cases such as CRIMINAL COURT become implausible. The group justification cannot diverge from the members such that no member shares the group's justification. Lackey's way out sacrifices the straightforward reading of divergence cases. Of course, she is aware of this and more than happy to explain these cases differently (cf. Lackey 2016: 354–7). I agree that this is an option, but I do not think it is the only one.

There are at least two alternatives one can aim for. First, one can argue against the claim that joint acceptance accounts have to endorse this sort of arbitrariness. Perhaps they can find a way to stop believing at will for groups. For instance, they can adopt Schmitt's (1994) idea from chartered groups having a structure, aim and ruleset in place that limits how group beliefs can be formed while still accepting a form of joint acceptance condition for group states.⁶ However, this is not the strategy I chose. My solution is different. I now proceed by developing an alternative that keeps JAA in the game while fully conceding the arbitrariness in joint acceptance. In doing so I can grant Lackey as much as possible, while still rejecting her conclusion. In itself this is not an argument against Lackey's own account of group justification. Whether Lackey's view or a form of JAA is correct will depend on whether divergence cases exist. I merely argue that IGNORING EVIDENCE and FABRICATING EVIDENCE cannot decide in favour of Lackey's view.

3. The Containment Strategy

Lackey presents IMEP as resulting in a dilemma: either one accepts that reasons can be arbitrary and therefore groups can manipulate their justification at will, or one rejects JAA. However, there is a third option: I propose that the proponent of JAA can bite the bullet on the arbitrariness of reasons, but deny that groups thereby can manipulate their *ultima facie* justification at will. Whereas Lackey attempts to get rid of the arbitrariness, I suggest that we merely contain its effect on the epistemic propriety of group states. The idea is to introduce conditions that set up boundaries of potential justification that can be achieved by a group's reasons. Call this the *containment strategy*. The aim is to contain the epistemic badness of arbitrarily formed reasons, without deferring to the epistemic status of the group members. To emphasise: my strategy fully endorses that joint acceptance accounts of group justification leads to arbitrarily formed reasons and hence *prima facie* justification at will. But I deny that this leads to implausible results of *ultima facie* justification because I can contain the arbitrariness. As long as I end up with the right results for *ultima facie* justification I do not think the unusual verdict for *prima facie* justifications should worry me.

The strategy is not unique to groups. In fact, Lackey herself accepts a form of the containment strategy for individuals, though it occurs under a different guise: normative defeat.⁷ Here is Lackey explaining normative defeat:

A normative defeater is a doubt or belief that S ought to have, and indicates that S's belief that p is either false (that is, rebutting) or unreliably formed or sustained

⁶Thank you to a reviewer for emphasising this option.

⁷Other proponents of normative defeat include Harman (1980), Kornblith (1983), DeRose (2000), Meeker (2004), Reed (2006), or more recently Goldberg (2016). Normative defeat only matches the containment part of my strategy, not the acceptance of arbitrary reasons as *prima facie* justification.

(that is, undercutting). Defeaters in this sense function by virtue of being doubts or beliefs that S should have (whether or not S does have them), given the presence of certain available evidence. (Lackey 2016: 366)⁸

As she tells us, normative defeat paradigmatically refers to evidence that one should have, regardless of whether one actually has the evidence. The classic example is due to Kornblith (1983), who invites us to imagine a physicist who regularly tunes out his colleagues' comments whenever they provide him counterevidence to his pet hypothesis. Kornblith convincingly argues that even if the physicist's current evidence justifies his belief in his pet hypothesis, this belief is overall not justified. There is evidence the physicist should have had. He should have listened to the counterevidence his colleagues told him about. Clearly, the physicist is doing something epistemically improper, and if he had been a virtuous physicist he would have heard the counterevidence and thereby lost his justification. The physicist's belief is normatively defeated by evidence he should have had. His epistemic impropriety is *contained* by the effects of the evidence he should have had: the normative defeater stops the improperly formed evidence basis from providing *ultima facie* justification. And it can do so without the physicist having any inclination of the normative defeater existing at all.

Normative defeat is not completely uncontroversial. Feldman (2000), for instance, holds that the single epistemic norm we have requires only a proportioning of one's belief to one's possessed evidence. There is no place for evidence one should have had, according to Feldman. Be that as it may, given that Lackey accepts normative defeat I take it as common ground available to defend JAA.

Kornblith's physicist case has clear parallels to the case of IGNORING EVIDENCE. It seems rather natural to me to describe IGNORING EVIDENCE as a case of Philip Morris not possessing evidence it should have had. Given that the members of Philip Morris have a reason to put warning labels on cigarette packs insofar as they believe that smoking is unhealthy, the group consisting of these members should have this reason as well. But the group does not have this reason available. This is explicitly highlighted by Lackey:

Consider IGNORING EVIDENCE: given that all of the evidence showing that smoking is dangerous is not available to the group because of the members' refusal to jointly accept it, none of it is part of the justificatory basis of the group's belief. (Lackey 2016: 352)

Here she is talking about evidence that is not available to the group because of the members' actions. Similarly, the physicist does not have the proper evidence available, because of his refusal to engage with it properly. None of the colleagues' counterevidence is part of the justificatory basis of the physicist's belief. Considering only the physicist's actual evidence it looks as if he is justified. But he is not. His justification is normatively defeated. Similarly, considering only the group's actual reasons that are formed by joint acceptance it looks as if the group is justified. And Lackey is right that in this case Philip Morris being justified sounds absurd. However, this can be explained by normative defeat just as the physicist case was. The group's justification is normatively defeated. *Ultima facie* justification is not merely a product of reasons the group has, but also of reasons the group should have. Hence, there is no need to abandon JAA just by raising IGNORING EVIDENCE.

⁸Lackey relies on normative defeat in other work as well. See for instance Lackey (2005).

This line of reasoning relies on the claim that beliefs of the group members can have a particular relation to normative defeaters for the group. In some cases we can gauge which normative defeaters apply to a group based on group members and their beliefs. Because the members of Philip Morris believe that smoking is unhealthy, we have reason to believe that Philip Morris has a normative defeater – the group lacks evidence it should have. This might be surprising, given that a JAA holds that justification of individual members does not straightforwardly transmit to the group level. Group justification has to be something over and above member justification. Therefore, I need to further unpack this claim: It is not mere belief of members that is a guide to normative defeaters for the group, rather it is *ultima facie* justified belief. Belief that is both *prima facie* justified and not normatively defeated. It is this latter component that matters for my claim. Whereas neither *prima facie* nor *ultima facie* justification of members transfers to the group level, the epistemic standards that make up normative defeat apply to individuals and groups equally.⁹ The same evaluative standards that determine what evidence an individual ought to have also apply to the group and determine what evidence the group ought to have. Hence, if we know that a particular belief of an individual member satisfies all epistemic standards we can infer whether the group's belief satisfies those standards as well. And in the case of Philip Morris it seems obvious that the group belief does not satisfy the same standards with regard to normative defeat. This is compatible with a JAA, because *ultima facie* justification of the group still is over and above *ultima facie* justification of individual members. The epistemic status of individual members itself does not directly impact the group justification. It is only our guide to epistemic standards that apply equally to individuals and groups, and these standards destroy the group's *ultima facie* justification via normative defeat.¹⁰ In contrast, Lackey's own view posits that the group's *prima facie* and *ultima facie* justification depend directly on a significant percentage of operative members justifiably believing a proposition (2016: 381). Her view is not compatible with JAA. I will come back to the relation of groups and their members with regard to standards for normative defeat in section 6, after presenting tools developed by Goldberg (2018) that will help me discuss the issue in more detail.

Even if these considerations are correct, JAA is not safe yet. Whereas *IGNORING EVIDENCE* looks structurally similar to Kornblith's physicist case, *FABRICATING EVIDENCE* is different. Here the group is creating evidence, not merely lacking evidence it should have had. How can the same strategy work? I take this case to be a little trickier, and it requires further discussion of the nature of normative defeat and epistemic propriety more generally. Fortunately, Sanford Goldberg provides much of the required ingredients to cook up a solution. I will first use his discussion to get a better understanding of what is going on in the *IGNORING EVIDENCE* case, and then explore how to explain *FABRICATING EVIDENCE* in a similar way.

4. Justification and General Expectations

Goldberg (2018) proposes that many types of evaluative assessments are hybrid assessments, built from core criteria and general expectations. The core criteria can either help or hurt the case for meeting the evaluative standard. General expectations on the other hand can only disqualify a subject from meeting the standard.

⁹I do not need the claim that they are exactly equal at this point, just that they are similar enough so that the standards for individuals can be used as basis of a heuristic allowing us to grasp the standards for groups. I suggest that they are equal in section 6 and provide a short argument there.

¹⁰Thank you to a reviewer who suggested to spell out this relation more explicitly.

There are 'hybrid' assessments involving evaluation along two distinct dimensions, where one of the dimensions includes assessment relevant considerations that can either support or bear against the claim that the subject *S* meets the standard φ , and the other of which includes assessment-relevant considerations whose only significance is as potentially bearing against (but not as providing positive support for) the claim that *S* meets φ . (Goldberg 2018: 55)

Goldberg suggests as a simple example a search committee trying to pick out new employees from a pool of candidates. The committee's job is to find out who is worthy of serious consideration for the job. Take this as the evaluative standard φ . The job description provides core criteria to meet this standard. Candidates need certain skills evidenced by education, previous job experience, and so on. Having a sufficiently good and fitting education and experience clearly helps to meet the standard. Lacking these would rule out a person as a potential employee. This is exactly what Goldberg predicts for core criteria: they can help, but they can also hurt the case for meeting φ . The committee can rank the candidates based on how well they do with regard to these core criteria. However, as Goldberg points out, these are not the only relevant factors that play a role for who is getting the job. Additionally, there are *general expectations* in place that can disqualify a candidate. Someone might be high-skilled, educated and experienced, but a terrible co-worker. Perhaps they keep talking loudly to themselves in shared work environments, dress inappropriately for the job, or keep harassing colleagues. Even the highest ranked candidate given the core criteria might be disqualified for being an unsuitable co-worker. These general expectations can disqualify, but they can never help a candidate who fails to satisfy φ . Being a great colleague cannot compensate for not having the ability to do the job properly. If one does not know how to write code in C#, then one is ruled out from the job as a C# coder, regardless of being a great person.

Goldberg provides several other examples that I cannot go into here (see Goldberg 2018: Ch. 2). In general, the core criteria are meant to capture the standard itself, or features that are good indicators for the standard being met. The general expectations need not indicate that the standard is met at all. They are background assumptions that assessors are entitled to take for granted. Crucially, these expectations are not arbitrary. I cannot simply decide that I expect an applicant to, say, like lemon tarts and that not liking lemon tarts disqualifies the applicant. Hence, Goldberg gives us a characterisation of these expectations. First, they must fit, or "make sense" (Goldberg 2018: 60) in the broader domain the assessment takes place and what the assessment is meant to accomplish. Obviously, constantly talking loudly to oneself when the job is located in a shared office environment is not something the committee looks for. Similarly, harassment is not something desirable in general, and for a work climate in particular. The domain, including relevant social actors, determine what makes sense as a general expectation. Moreover, these expectations can be changed through cultural changes. What makes for a good co-worker today might not be exactly the same as in the 1800s.

Second, Goldberg proposes that general expectations ought to be *legitimate*. Legitimacy here relates to our responses to potential challenges of these expectations. Some expectations are rightly taken for granted. Someone might complain that the job description did not list that people are expected not to harass co-workers and that this expectation is arbitrary. But clearly, the committee could respond by providing reasons for this general expectation – reasons that would be generally accepted in the relevant social environment. Goldberg admits that it is not always clear whether an expectation is legitimate, but I take it that there are at least clear cases in which one is happy to accept that the expectation is legitimate. That is all I require.

After having discussed Goldberg's general picture I can continue with the case of epistemic evaluations. These work in the same fashion as the evaluation of job candidates. They are also hybrid between core criteria and general expectations. And again, the core criteria can either help or hurt, whereas the general expectations can only disqualify. This is where normative defeat comes into the picture. A subject's justification is normatively defeated when some general expectation of evidence acquisition is not met. The subject lacks evidence that they are generally expected to have in their situation. It is *evidence they should have had*. It is important for my purpose that these general expectations hurt the case for meeting the epistemic standard. The containment strategy aims to stop the epistemic badness of arbitrary evidence from leaking into ultima facie justification. Normative defeat is supposed to do the job. And it can do so by us having general expectations on the evidential state a subject should have. In *IGNORING EVIDENCE* Philip Morris has an evidential state that does not fit our general expectations for the situation. We expect Philip Morris to possess evidence for the effects of smoking on health. Moreover, we expect Philip Morris to therefore have a reason to stick warning labels on packs of cigarettes. Even if Philip Morris proportions its belief correctly to the actually held group reasons, Philip Morris does not meet the standard due to violating these general expectations. That is, Philip Morris might satisfy the core criteria, but that is not enough to overall satisfy the epistemic standard. Just like the high-skilled, rude applicant will be ruled out, even though they meet the core criteria. The general expectations stop a subject from being more justified than it ought to be by introducing a ceiling on ultima facie justification. Goldberg formulates this *epistemic ceiling principle (ECP)* already in an earlier work on normative defeat:

ECP: Given a subject S's belief that p based at time t on evidence E, if there was evidence S should have had (E_{SHH}) [*Evidence S should have had*] at t, then (i) S's belief that p at t can be no more justified than it is on E, and (ii) S's belief that p can be no more justified at t than it would have been if S had acquired E_{SHH} at t (and updated in a rational fashion). (Goldberg 2016: 462, italics added)

Here the hybrid nature is made obvious in virtue of the two distinct conditions: (i) captures a ceiling that is a product of the core criteria. (i) relates justification to the evidence a subject possesses. A subject cannot be more justified than the core criteria (the possessed evidence) allows it to be. This follows from the general view on hybrid assessments, because no factor besides the core criteria can help meet the standard in question. (ii) creates the upper limit included by general expectations of the evidence a subject ought to possess, regardless of whether the subject actually has this evidence. This captures normative defeat. ECP explains why Philip Morris is not justified in *IGNORING EVIDENCE*. The evidence jointly accepted determines the upper limit of justification based on (i) and the evidence the group should have had determines the upper limit of justification based on (ii). Philip Morris cannot be more justified than the lower bound of (i) and (ii). In this case (ii) makes sure that Philip Morris cannot be justified at all, hence the group is not justified to believe that it should not put warning labels on cigarette boxes.¹¹

¹¹One might object here, that this is not quite the same as in *IGNORING EVIDENCE*. The question is not about being unjustified regarding the belief that the group should *not put* warning labels on the boxes. It is about being justified to *put* these labels on the boxes. As a first response this cannot be right, because Lackey is clear that she is talking about doxastic, not propositional justification (Lackey 2016: 342–3). And justification for a belief the group does not hold cannot be doxastic. As a second response I point to the

One might question at this point whether that is needlessly complicated for IGNORING EVIDENCE. Why not just argue that the core criteria are not met?¹² The answer to this question depends on how exactly we conceive of the core criteria. These depend on the background theory we accept. Goldberg makes this explicit when he states that “there are a variety of different theories each aiming to articulate such standards: reliabilism, evidentialism, coherentism, phenomenal conservatism, and so forth” (Goldberg 2018: 69). Goldberg himself opts for a version of process reliabilism. I also commit to process reliabilism as the background for understanding core criteria. Like every account of justification on offer reliabilism has its fair share of challenges, but these will not be solved in this paper. I merely want to highlight the background commitment to reliabilism in both Goldberg and my understanding of core criteria. Given this understanding I can explain how the core criteria can be met in IGNORING EVIDENCE, and hence why the general expectations need to do the job of removing *ultima facie* justification. Jumping back on Schmitt’s (1994) idea of a chartered group we can identify an office and a charter for Philip Morris. To remind us: a group’s charter captures rules, norms and standards that organise the group and its members. These rules need not be written down, but the group members need to be aware of these rules on some level. Philip Morris’s charter regulates how decisions are made in the company, and how new economic targets are determined. The office is the group’s aim and purpose. For simplicity we can just assume it is the aim of Philip Morris to maximise profits. I am happy to concede that these toy versions of charter and office are much simpler than real world cases, but they will do a good enough job for my purpose. Together charter and office provide a background under which the group forms group states. I take charter and office to leave plenty of room for arbitrariness in joint acceptance. However, they provide a way to individuate the group’s belief formation. The group’s belief-forming process can be individuated as joint acceptance by the operative members following the group’s office and charter. Now the reliabilist take on core criteria just has to take this belief-forming process and look at its reliability. And it is possible to tell the tale of IGNORING EVIDENCE in a way that makes Philip Morris reliable enough to satisfy the threshold required for the core criteria in the reliabilist version. This is merely a result of combining the joint acceptance idea with process reliabilism. Some might feel reluctant to accept this result, but as long as the verdicts for *ultima facie* justification match our intuitions I do not see a convincing reason to abandon this framework. Hence, the importance of

related idea of propositions that a subject *should have known* (cf. Goldberg 2017). Some evidence that we expect subjects to have comes in the form of beliefs, or even knowledge. For instance, suppose you are part of a protest wielding a huge poster stating ‘stop global warming’. It seems natural to assume that there are general expectations about you knowing what ‘global warming’ means or refers to. I expect you to know about global warming. If you do not, there is a clear sense in which you are blameworthy. You should have known!

Similarly, some general expectations of knowledge go with working in a particular domain. If you work as a geologist, one expects you to have the relevant knowledge. If you do not, you might be rightfully blamed. And Philip Morris? Would we not generally expect an organisation in the tobacco domain to know about the effects of cigarettes? I certainly do. Someone who produces a product *should know* about the effects of said product. IGNORING EVIDENCE shows Philip Morris’ failure on this axis of general expectations. Philip Morris should have known what the studies about the effects of smoking state. Philip Morris should have a reason to believe that it should put warning labels on cigarette boxes. Hence, we can blame Philip Morris for not having the right propositional justification by considering the general expectations of what Philip Morris should know.

¹²Thank you to a reviewer for raising this issue.

the general expectations to explain why the group is not *ultima facie* justified in IGNORING EVIDENCE. These give us the right verdicts for *ultima facie* justification.

General expectations are also the key to deal with FABRICATING EVIDENCE. Here again the arbitrariness of group reasons creates a problem. If the group can form reasons at will, then the group can create justification for any belief of their choosing out of thin air. Philip Morris decides to jointly accept that particular scientists are liars and thereby gets a reason not to believe their studies. Just looking at the reasons that Philip Morris possess this belief seems justified. But clearly, it is not.

I previously wrote that this case is different to Kornblith's scientist scenario. It is different because the scientist merely ignored evidence, but did not create its own evidence at will. However, if we conceive of these cases as more general violations of general expectations they turn out to be closely related. Normative defeat functions not merely through a lack of particular evidence. Rather, normative defeat is based on the evidence actually possessed coming apart from the evidence one ought to have. This can be because the subject lacks evidence, but it may also be because the subject has evidence it should not have.¹³ In FABRICATING EVIDENCE Philip Morris clearly has evidence that it should not have in addition to lacking evidence it should have. The group's evidential profile is not even close to what one expects it to be, given the situation Philip Morris is in.

Let me illustrate this with a modified version of CRIMINAL COURT:

FAILING JURY: A defendant is on trial for the crime of careless driving. The prosecution adduces evidence from the police report, as well as eyewitnesses testifying in court that the defendant was indeed driving the van that hit the victim. But the jury finds the evidence not beyond a reasonable doubt, and hence insufficient to validate criminal conviction. All the members of the jury have hearsay evidence from a reliable source that the defendant caused the accident. But the judge instructs the jurors to ignore this evidence as it fails to meet the conditions for being ruled admissible in a criminal court. However, the jury ignores the judge, even though it is functioning in its social office, and bases its belief that the defendant is guilty on the hearsay evidence.

In FAILING JURY the group has a justified belief that the defendant is guilty, according to JAA. But it seems to me that the jury should not have that justified belief. The jury uses evidence it should not have and is therefore blameworthy. This is a result of general expectations towards the jury. A jury is expected to ignore evidence that the judge orders them to dismiss. Here this expectation is violated. And just like in FABRICATING EVIDENCE, the group ends up with a reason it should not have. The group, functioning in its social office, should not be justified in believing that the defendant is guilty, regardless of whether the group based that belief properly on the group's evidence. The jury's job can be assessed on two levels: they satisfy (i) the core criteria of basing their judgment on their evidence. However, they do not meet (ii) the general expectations of having a particular evidential profile on basis of which they make their judgments. It is part of these general expectations to not include evidence that they should not have. And they cannot be more justified than they would be without this piece of evidence that they should not have.

¹³This marks a divergence of my take on general expectations from Goldberg's. Goldberg only addresses cases in which a subject should have more evidence, he does not discuss cases in which someone has evidence they should not have. Thank you to a reviewer for emphasising my departure from Goldberg's view here.

One might object here (as Lackey does) that the jury case is not about epistemic justification, but rather some practical concern, or about an idiosyncratic fact of the social office in place.¹⁴ But even so, what the case shows is that there are hybrid evaluations that include as part of their general expectations a comparison between what the actual basis for a judgment is, and what the basis should be. If this structure exists for assessing whether the jury meets some judicial standard, then the structure might also exist for epistemic standards. I do not need to rely on the jury case to be read epistemically. Take it as a mere proof of concept. I now show that expectations with a similar structure can do the job required by using them to explain FABRICATING EVIDENCE.

I propose to deal with FABRICATING EVIDENCE by suggesting such a structure for epistemic standards. Evaluating epistemic propriety is a hybrid assessment based on (i) the core criteria of proportioning one's belief to one's evidence, and (ii) the general expectation of having a particular evidential profile in a given situation. We can use this to formulate a modified ECP accordingly:

ECP*: Given a subject S's belief that *p* based at time *t* on evidence *E*, if the actual evidence that S has at *t* differs from the evidence S should have at *t*, then (i) S's belief that *p* at *t* can be no more justified than it is on *E*, and (ii) S's belief that *p* can be no more justified at *t* than it would have been if S had the evidence S should have had at *t* (and updated in a rational fashion).

ECP* gives us the correct verdicts for both IGNORING EVIDENCE and FABRICATING EVIDENCE. In both cases the overall justification Philip Morris has cannot be better than the justification a subject with the proper evidence would have had. Regardless of whether Philip Morris has too much, or not enough evidence, compared with the evidence it should have had. Given ECP* arbitrariness of reasons cannot leak into arbitrariness of justification. It does not matter how many additional reasons the members of Philip Morris jointly accept, the group's justification cannot break out of the boundaries set by the evidence the group ought to have. The badness of their reason-formation is contained by the normative defeat generated through general expectations. If ECP* holds, then JAA can explain Lackey's challenges sufficiently. In both cases of IGNORING EVIDENCE and FABRICATING EVIDENCE Philip Morris is not justified. This is exactly what the result ought to be. There is no need to abandon JAA.

Unfortunately, ECP* seems to generate a problem. A reviewer raises concerns about ECP* in cases of a person being epistemically overzealous in a competent fashion. Students in a class might not be expected to do additional research, but suppose an especially talented student does more than expected and handles the additional evidence from her individual research with care and competence. ECP* seems to claim that the student is no better off epistemically than all the other students who did not do that research, because no one can be better justified than the evidence one should have at *t* determines. But how can more carefully and competently gained evidence not result in better *ultima facie* justification? Does ECP* give a wrong verdict in these cases?

Let me take a closer look at the student who does competent additional research on her own, even though she is expected not to do so. One way to read the case together with ECP* is to argue that she has more evidence than expected, but cannot be more *ultima facie* justified than someone with less evidence is, because she should have less evidence. This result is unpalatable. In this reading we set the expectations at the

¹⁴As a reviewer worries.

start of the class at t_1 , and then the student gathers additional evidence, leading to her having more evidence than she should have at t_2 . However, there is an alternative reading available. Expectations are not set in stone. General expectations are determined by considering an epistemic subject in their particular position and situation. The problematic reading of the case assumes that the expectations are determined at t_1 in the situation at the start of the class and the position the student is in at that moment. However, I suggest that the student's additional research does not merely change the evidence she has, but also the position and situation she is in. That is, the research changes the general expectations of the student's epistemic state. The evidence a student should have after additional research differs from the evidence a student should have without the additional research. ECP* is therefore compatible with the overzealous, competent student having better *ultima facie* justification than the other students in class. It is compatible, because the evidence S should have had at t_2 is different from the evidence that she should have at t_1 and also different from the evidence the less zealous students should have at t_2 . The upper limit constituted by (ii) is higher for a student in a position that provides better access to potential evidence. We just expect more from the person who did more research.¹⁵ And that seems to be exactly the verdict we want.

5. Core Criteria

I have been talking mainly about the general expectations that explain normative defeat and its relevance for IGNORING EVIDENCE and FABRICATING EVIDENCE. But what about the core criteria? Epistemic assessments are hybrid between core criteria and general expectations. What if the core criteria are already a problem for JAA? At first sight this seems to be a natural response. Goldberg (2018) takes the core criteria to involve reliability. Lackey (2016) thinks the problem of JAA is their lack of grip on the world. Should Lackey therefore respond that JAA fails the core criteria and therefore all these considerations of general expectations cannot save JAA?

There are two different avenues of response. One is a flat-footed internalist one. One might be generally in agreement with Goldberg's hybrid picture of epistemic evaluation, but disagree about the nature of the core criteria. Perhaps all that matters for the core criteria is that one believes in accord with one's internalist evidence, regardless of how that evidence relates to the world.

My preferable, second response is closer to the reliabilist's position. Classic process reliabilism holds that a belief that p is justified if it results from a belief-forming process that generates a high enough ratio of true beliefs. And in this standard picture belief-forming process refers to some cognitive operation with which a "cognizer deals with his environmental input" (Goldman 2008: 340). It is not clear how this ought to be applied to groups. In particular, how belief-forming processes ought to be individuated for groups. If all group beliefs are formed by joint acceptance, they all seem to be generated by the same process. The generality problem (Conee and Feldman 1998) seems to be a particularly pressing challenge here. Moreover, for individuals there is an explanation for reliability available. Vision has evolved to give us access to the world in a particular way that produces reliably true beliefs. Of course, not always, but reliably

¹⁵Importantly, ECP* still states that we do expect the student to not have more evidence than the additional research warrants. What (ii) of ECP* targets are situations in which an epistemic subject has evidence that should not be in the subject's possession given the position the subject is in. Of course, in judging the epistemic position of other people we might make mistakes and misjudge the position someone is in. We might have false expectations because we are not fully aware of the situation. But in these cases the problem is not with ECP*, but with our judgment.

enough. If JAA is right, groups form beliefs at will. Why should these beliefs be formed reliably? There is nothing in the process that secures reliability. There is a flat-footed reliabilist response available claiming that it does not matter what the process looks like. All that matters is that it produces true beliefs often enough. If, for some groups, their joint acceptance generates true beliefs often enough, then these groups satisfy the reliabilist criteria.

However, this appears to fall victim to the ability intuition: for knowledge pure reliability is not enough. The reliability has to be a manifestation of a competence (see for instance Sosa 2007, 2015; Greco 2009; Pritchard 2012). Plausibly, the same idea applies to justification. Reliability that is not a result of one's competence does not seem enough for justification. Consider Pritchard's Temp case:

TEMP: Temp forms his beliefs about the temperature in the room by consulting a thermometer. His beliefs, so formed, are highly reliable, in that any belief he forms on this basis will always be correct. Moreover, he has no reason for thinking that there is anything amiss with this thermometer. But the thermometer is in fact broken, and is fluctuating randomly within a given range. Unbeknownst to Temp, there is an agent hidden in the room who is in control of the thermostat whose job it is to ensure that every time Temp consults the thermometer the 'reading' on the thermometer corresponds to the temperature in the room.

Temp is reliable, but does not appear to be justified. The truth of Temp's beliefs does not latch on to the world in the right way. Hence, brute reliability does not seem to be enough. Fortunately, JAA is not committed to joint acceptances that cannot be manifestations of competences. Friends of JAA should respond that even though joint acceptances are not relevantly constrained by the competences of individuals, they can be constrained by the group. And if they are, they manifest competences on the group level (cf. Kallestrup 2016). These group competences are based on the structure of the group. One way to understand the group's structure is to come back to chartered groups as mentioned earlier. The structure of these groups can be understood based on the aims and goals that constitute the group (the office) and the rules, norms and standards that organise the group's and group member's conduct (the charter). This is not the only way to understand the structure of a group¹⁶ but is instructive for my purpose. Consider again Philip Morris. Its office includes the practical goal of selling packs of cigarettes and creating revenue. Its charter includes the organisational structure, how they decide on actions as a group, who is an operative member, etc. For simplicity assume that the charter includes that decisions are formed by unanimous voting of the executive board. Together charter and office provide a background under which the group forms group states. The idea here is that groups do not start to form group acceptances for no reason. Rather, groups build around joint intention, the fulfilment of which requires further group acceptances. The group exists to fulfil its office. Moreover, these joint intentions lead to organisational structures to support the fulfilment of these intentions. The charter captures this structure and is meant to solve the organisational question of how the members work together to achieve the group's office. As already pointed out earlier, I can use this idea of a group's structure in terms of charter and office to individuate the process that ought to be reliable. A group's attitude-formation process type is individuated in terms of the joint acceptance of an attitude by the members of a group following a certain charter and office. Charter and office are group level features, therefore this does not require me to refer to states of individual members. If a

¹⁶For a different approach to groups as structured see Ritchie (2013).

group following a charter and office generates a sufficient ratio of true beliefs/acceptances, then the group forms these states reliably, even if the group's office does not include an epistemic goal at all. Or in other words: the combination of charter and office constitute the group's competences with regard to forming beliefs. All beliefs that are formed by joint acceptance following the group's charter and office are manifestations of the group's competence. They therefore fulfil the ability intuition.

To come back to the example, I can now ask whether Philip Morris with its particular office and charter accepts propositions in a way that results in true beliefs often enough to pass a reliability threshold. That is: is Philip Morris' belief-forming process reliable? Here I simply plug in the flat-footed response from before: how many of the beliefs the operative members of Philip Morris jointly accept under the office and charter of the company turn out to be true? If enough of them are true, then the group counts as reliable and therefore *prima facie* justified. The ability intuition is satisfied, because Philip Morris formed these states in virtue of its charter and office, and thereby manifested its competence. If Philip Morris is reliable, then the reliability is explained by pointing to its charter and office. A different charter and office would have led to different jointly accepted states.

Even if reliability is established one might be worried at the cost of endorsing process reliabilism.¹⁷ I am committed to the following result: Suppose that a group G is 99% accurate in its acceptances but once a month they decide to accept a proposition just for the fun of it. Do I want to say that this just for fun acceptance is *prima facie* justified because it's generated by a reliable process? My answer here has to be a clear 'yes', if the acceptance just for fun is compatible with the group's office and charter. This is the arbitrariness I endorse for JAA. If I want to defend these accounts without giving up arbitrariness I have to answer 'yes' here. However, the containment strategy will stop this arbitrariness from leaking into *ultima facie* justification. And I am now in a position to look back to the hybrid nature of epistemic evaluation to further explain the containment strategy for FABRICATING EVIDENCE.

Remember that I am looking for core criteria and general expectations. I considered the apparent problem that groups have a difficult time satisfying the core criteria, so that the containment strategy focusing on general expectations enters the picture too late. I have now provided a way to think about the core criteria that leads to criteria that groups can satisfy. Groups can be evaluated according to their production of states in virtue of their charter and office. Thereby a group can satisfy the core criteria proposed by Goldberg. Let me illustrate the final picture with FABRICATING EVIDENCE:

FABRICATING EVIDENCE: Philip Morris is one of the largest tobacco companies in the world, and each of its operative members is individually aware of the massive amounts of scientific evidence revealing not only the addictiveness of smoking but also the links it has with lung cancer and heart disease. Entirely because of what is at stake financially and legally, however, each of these members decides to jointly accept that all of the scientists working on the relationship between smoking and health problems are liars. Given this, they also jointly accept that the duplicity of the scientists gives Philip Morris a reason to believe that the results of the studies showing a connection between smoking and lung cancer and heart disease are unreliable. (Lackey 2016: 351)

In an epistemic evaluation I ask first whether the core criteria are satisfied. Plausibly, that is the case in FABRICATING EVIDENCE. Most of Philip Morris beliefs/acceptances

¹⁷Thank you to a reviewer for pressing me on this point.

are true, evidenced by their success in fulfilling their office and increasing their revenue. It is unlikely that Philip Morris is so successful without getting things about the world right. For instance, the group surely knows what products the customers tend to buy, what the best way to distribute the product is, etc.. Hence, Philip Morris is generally reliable in virtue of its charter and office. Moreover, in this particular case the group produced evidence by joint acceptance in line with charter and office. Accordingly, Philip Morris' act of proportioning its belief to the evidence is enough for prima facie justification in this particular case as well. The end result is that the belief that the studies on the effects of smoking are unreliable is prima facie justified and the core criteria are satisfied.

However, Philip Morris does not satisfy the general expectations, because in this particular instance Philip Morris' evidence diverges too far from the evidence Philip Morris should have had. Philip Morris' belief that these studies are unreliable can be no more justified than it would have been if Philip Morris had the evidence it should have had, where the evidence it should have had is determined by considering what evidence an individual in the group's position is expected to have. In FABRICATING EVIDENCE Philip Morris is prima facie justified due to fulfilling the core criteria, but is not ultima facie justified due to failing to meet the general expectations. This is the right result, and it does not require sacrificing JAA.¹⁸

6. An objection and the origin of group expectations

In the last sections I developed the containment strategy and showed how to employ it to defend JAA against Lackey's challenging cases. In doing so I relied on Goldberg's idea of general expectations and their role in epistemic assessments. A natural worry to the containment strategy is that it might appear to be an ad hoc manoeuvre. I do not think it is. Expectations are part of epistemic evaluation in general. They exist for individuals, so it should not surprise anyone to find them again at the group level. If Goldberg is right, it is just a general feature of epistemic assessments that they are hybrid. Of course, at this point one might concede that general expectations are in place for individuals, but we need an independent reason for them being in place for group agents. After all, there is a story available why these expectations are in place for individuals. Goldberg argues that practical concerns establish these sorts of expectations. Our social practices of information-exchange and practical engagement with each other demand that all participants follow certain epistemic norms. Participation in these practices would be practically irrational without an entitlement to these expectations. I cannot rationally take part if I do not take the other participants to follow these norms of epistemic propriety. Is this kind of argument also applicable to groups? I am not so sure. On one hand, group actors can be part of our information-exchange practices, if we allow groups to testify (see for instance Tollefsen 2007; Fricker 2012; Lackey 2018). Hence, group actors might fit into the same story about practical concerns. On the other hand, it is easy to imagine that groups played no role in how these norms of epistemic propriety actually came about. They might pre-date anything like group testimony. However, this does not

¹⁸A reviewer worries that this commits me to say that if our general expectations for Philip Morris were different, Philip Morris could get ultima facie justification from arbitrary reasons. I am happy to bite the bullet, because this is an 'if' that is unlikely to be actualised. Remember that Goldberg argues that practical concerns establish these sorts of expectations. Our social practices of information-exchange and practical engagement with each other demand that all participants follow certain epistemic norms. I take it to be very unlikely that our practical needs allow for expectations that let the arbitrariness leak into ultima facie justification. Such leakage would result in unreliable testimonial practices.

seem to be a problem. All that my strategy requires is that we apply the same standards of epistemic evaluation to groups and individuals. It is compatible with these standards being formed in response to practical problems involving only individuals. Once the standard for individuals is in place it can be repurposed to also evaluate the epistemic status of groups. One does not need new standards for epistemic propriety of groups. One merely needs to be able to see groups as possible epistemic agents to apply the standard that is already in place.

Reapplication of individual epistemic standards to groups is not only economic, but provides an explanation of what evidential profiles we compare to assess the justification ceiling. ECP* claims that S's belief that *p* can be no more justified at *t* than it would have been if S had the evidence S should have had at *t*. To do that, we need a grip on what evidence S should have had. For individuals, this is not especially difficult, because we know the belief-forming processes of individuals and how they can be used in particular situations. We have an intuitive grasp of what is expected for an individual in a particular position at a particular time. For groups, this seems more problematic. If JAA is right, then groups form beliefs (or joint acceptances) in a way that has little in common with belief-formation in individuals. Hence, it might be unclear what evidential profile an agent with unfamiliar belief-formation ought to have. Even worse, according to JAA groups form their states by sheer willingness to let an attitude that *p* stand as a group attitude. Why should we have general expectations about evidence for these states? They are arbitrary after all, and general expectations of wholly arbitrary states seem misplaced. However, if the epistemic evaluation of group states is merely a reapplication of the standards we use for individuals, this question does not arise. We have general expectations for individuals that include having acquired adequate evidence for particular situations. We then reapply these expectations to group agents and expect that a group has acquired evidence that would be adequate for an individual in the group's situation. The general expectations are determined by considerations of individuals. Therefore, there is no need to develop new general expectations for groups. I expect this to be a lot messier than I present it here. Even for individuals it is not always obvious what general expectations we have. There are clear cases. We know that the physicist is expected to listen to results of the experiments of other physicists. But there are a lot more cases in which we are on the fence on whether someone should have known a particular piece of information. Even more so, it is unclear what individual we imagine in place of the group when we determine the general expectations towards the group. Do we imagine an expert, or a novice? Do we imagine someone who is an 'average' member? I do not have a final answer to these questions. My tentative answer is that we use heuristics that draw on knowledge of the group's structure and practical considerations to imagine a single person with the resources and available tools of the complete group, and then ask what evidence we should expect such a person to have. I am happy to concede that this is not a fully satisfying answer and will require additional work in the future.

One might object, that reapplying the standards from individuals to groups is overall not motivated. The belief formation process of groups is very different from individual belief-formation. So why should the same standard of general expectations apply? I take it that here the practical considerations kick back in. Evaluating groups based on our standards for individuals allows us to include groups in our general practices of information-exchange. If I want to include groups in my information-exchange practices the evaluative standards should align. After all, I need to be able to relate information given by groups to information given by individuals. In particular, I can turn a piece of information taken from a group into a piece of information that I testify as an individual. If the same piece of information can be used in both, they had better

follow the same standards of propriety. Otherwise this simple transformation would endanger the information-exchange practices between individuals. If the epistemic standard for groups were different than for individuals, we would always risk one contaminating the other. This risk would remove the entitlement that we have towards the epistemic practices of other individuals. Hence, if the standard for groups were different than for individuals, our practices between individuals would become irrational. I take it that they are not irrational, hence the standards for groups have to be the same as they are for individuals.

7. Conclusion

In this paper I provided an attempt to defend JAA against proposed counterexamples by Jennifer Lackey. I argued that if we take Goldberg's suggestion of thinking of epistemic assessment in terms of a hybrid evaluation we can find a way to explain Lackey's examples without abandoning JAA. Epistemic evaluations are based on both core criteria and general expectations. In Lackey's cases the core criteria are satisfied, but the general expectations are violated. I suggest that friends of JAA take this hybrid nature of evaluation on board. In doing so they can keep the benefits of JAA without clashing with our intuitive judgments about manipulating evidence.¹⁹

References

- Bird A. (2010). 'Social Knowing: The Social Sense of 'Scientific Knowledge'.' *Philosophical Perspectives* 24(1), 23–56.
- Conee E. and Feldman R. (1998). 'The Generality Problem for Reliabilism.' *Philosophical Studies* 89, 1–29.
- DeRose K. (2000). 'Ought We to Follow Our Evidence?' *Philosophy and Phenomenological Research* 60(3), 697–706.
- Feldman R. (2000). 'The Ethics of Belief.' *Philosophy and Phenomenological Research* 60(3), 667–95.
- Fricker M. (2012). 'Group Testimony? The Making of A Collective Good Informant.' *Philosophy and Phenomenological Research* 84(2), 249–67.
- Gilbert M. (1989). *On Social Facts*. London: Routledge.
- Gilbert M. (1990). 'Walking Together: A Paradigmatic Social Phenomenon.' *Midwest Studies in Philosophy* 15(1), 1–14.
- Goldberg S. (2016). 'On the Epistemic Significance of Evidence You Should Have Had.' *Episteme* 13(4), 449–70.
- Goldberg S. (2017). 'Should Have Known.' *Synthese* 194(8), 2863–94.
- Goldberg S. (2018). *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford: Oxford University Press.
- Goldman A. (2008). 'What Is Justified Belief?' In E. Sosa, J. Kim, J. Fantl and M. McGrath (eds), *Epistemology: An Anthology*, 2nd edition, pp. 333–47. Oxford: Blackwell.
- Greco J. (2009). 'Knowledge and Success from Ability.' *Philosophical Studies* 142(1), 17–26.
- Hakli R. (2011). 'On Dialectical Justification of Group Beliefs.' In H.B. Schmid, D. Sirtes and M. Weber (eds), *Collective Epistemology*, pp. 119–53. Frankfurt: Ontos.
- Harman G. (1980). 'Reasoning and Evidence One Does Not Possess.' *Midwest Studies in Philosophy* 5(1), 163–82.
- Kallestrup J. (2016). 'Group Virtue Epistemology.' *Synthese*. doi: 10.1007/s11229-016-1225-7.
- Kornblith H. (1983). 'Justified Belief and Epistemically Responsible Action.' *Philosophical Review* 92(1), 33–48.
- Lackey J. (2005). 'Memory as a Generative Epistemic Source.' *Philosophy and Phenomenological Research* 70(3), 636–58.

¹⁹Thank you to Miguel Hoeltje, Raphael van Riel, Thorsten Sander and anonymous reviewers for Episteme and Synthese for their comments. A first version of this paper has been presented at a research seminar at the University of Duisburg-Essen. Thank you to all attending students for their feedback.

- Lackey J. (2016). 'What is Justified Group Belief?' *Philosophical Review* **125**(3), 341–96.
- Lackey J. (2018). 'Group Assertion.' *Erkenntnis* **83**(1), 21–42.
- Meeker K. (2004). 'Justification and the Social Nature of Knowledge.' *Philosophy and Phenomenological Research* **69**(1), 156–72.
- Pritchard D. (2012). 'Anti-Luck Virtue Epistemology.' *Journal of Philosophy* **109**(3), 247–79.
- Reed B. (2006). 'Epistemic Circularity Squared? Skepticism about Common Sense.' *Philosophy and Phenomenological Research* **73**(1), 186–97.
- Ritchie K. (2013). 'What are Groups?' *Philosophical Studies* **166**, 257–72.
- Schmitt F.F. (1994). 'The Justification of Group Beliefs.' In F.F. Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*, pp. 257–87. Lanham, MD: Rowman & Littlefield.
- Sosa E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. 1. Oxford: Oxford University Press.
- Sosa E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Tollefsen D. (2007). Group Testimony. *Social Epistemology* **21**(3), 299–311.
- Tollefsen D. (2009). 'Wikipedia and the Epistemology of Testimony.' *Episteme* **6**(1), 8–24.
- Tuomela R. (2013). *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.

Lukas Schwengerer is currently a postdoctoral researcher at the University of Duisburg-Essen. He received his PhD at the University of Edinburgh in 2018 for work on a transparency account of self-knowledge. His research is located at the intersection of philosophy of mind and epistemology, with a particular interest in self-knowledge and how anti-individualistic approaches in the philosophy of mind impact epistemological question. His current focus includes epistemological questions in the context of social groups, and research on the transformation of our epistemic practices due to modern technology. His work has appeared in journals such as *Erkenntnis*, *Episteme* and *Review of Philosophy and Psychology*.