

An Anchored Joint Acceptance Account of Group Justification

Lukas Schwengerer (Schwengerer.Lukas@gmail.com)

Penultimate draft – please cite the published version

<https://onlinelibrary.wiley.com/doi/10.1111/theo.12460>

Abstract

When does a group justifiedly believe that p ? One answer to this question has been developed by Schmitt (1994) and Hakli (2011): when the group members jointly accept a reason for the belief. Call this the joint acceptance accounts of group justification. Their answer has great explanatory power, providing us a way to account for cases in which the group's justification can diverge from the justification individual members have. Unfortunately, Jennifer Lackey (2016; 2021) developed a powerful argument against joint acceptance accounts. She argues that these accounts lead to epistemically arbitrary reasons and therefore justification at will. Group justification loses the necessary connection to the world to be truth-conducive. In this paper I develop a new form of a joint acceptance account that can deal with Lackey's examples: the anchored joint acceptance account of group justification. I argue that properly understanding the role of epistemic expectations can help us form the best version of a joint acceptance account. While justification is only generated by joint acceptance of evidence, the evidential expectations towards a group are anchored in the group members. This anchoring guarantees that groups cannot manipulate their ultima facie justification illegitimately.

1. Introduction

When does a group justifiedly believe that p ? One answer to this question has been developed by Schmitt (1994) and Hakli (2011): when the group members jointly accept a reason for the belief. Call this the *joint acceptance accounts of group justification* (JAA). Their answer has great explanatory power, providing us a way to account for cases in which the group's justification can diverge from the justification individual members have. Schmitt points to juries that might have constraints on their evidence that the individual jury members do not have. For instance, as a jury they are not allowed to consider hearsay evidence. Hakli (2011, pp. 120-121) gives the example of a weather forecast group, in which no member believes a particular forecast, but the group accepts one collective forecast anyway.¹ JAA seems ideal to give an appropriate verdict in these cases.² Hence, they provide a fruitful starting point for group epistemology. Unfortunately, Jennifer Lackey (2016; 2021) developed a powerful

¹ For other examples see Tollefsen (2009, p. 12) and Bird (2010, pp. 34-35).

² Though Lackey (2016; 2021) provides alternative explanations for the intuitions in these cases.

argument against JAA. She argues that these accounts lead to epistemically arbitrary reasons and therefore justification at will. Group justification loses the necessary connection to the world to be truth-conducive. In this paper I develop a new form of a joint acceptance account that can deal with Lackey's examples: *the anchored joint acceptance account of group justification*. I argue that properly understanding the role of epistemic expectations (Goldberg, 2018) can help us form the best version of JAA. While justification is only generated by joint acceptance of evidence (or reasons³), the evidential expectations towards a group are anchored in the group members. This anchoring guarantees that groups cannot manipulate their *ultima facie* justification illegitimately, while still providing a good explanation for cases in which group justification diverges from the individual group member justification.

I made a first attempt at using epistemic expectations to defend JAA in previous work (Schwengerer, 2021), but that attempt is insufficient. I use Lackey's (2016; 2021) argument against JAA and my earlier response as the first stepping stones towards a new account of group justification. I start by quickly explaining Lackey's argument against JAA in part 2 and my 2021 response to the argument in part 3. I then show that the response has a noticeable gap that I close with an account of evidential expectations for groups in part 4. In Part 5 I use the results from part 4 and provide a new account of group justification. Finally, I revisit Lackey's argument and show how my new account is both similar and different from Lackey's own account of group justification in part 6.

2. Lackey's Argument from Illegitimately Manipulated Evidence

In developing her own account of group justification Jennifer Lackey argues against current competing views. I will only be addressing *joint acceptance accounts* (JAA) of group justification and Lackey's argument against JAA as this is the basis for my previous (2021) suggestion for JAA. Two versions of JAA are developed by Frederick Schmitt (1994) and Raul Hakli (2011). Lackey (2016; 2021) presents these as following:

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

³ I will use these interchangeably.

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, p. 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, p. 348)

Both versions of JAA understand justification in terms of reasons⁴ that a group has. Reasons can support a group's belief that p – they justify the belief. But many accounts of group justification understand justification in terms of reasons a group possesses. The unique feature of joint acceptance accounts is that these reasons are established by the group members jointly accepting a reason. If the group members jointly accept that r is a reason for the group, then this joint acceptance is sufficient for r being a reason for the group. On some version this also requires that the group members openly express⁵ their willingness to accept r as a reason for the group, or conditions of common knowledge, but for mine and Lackey's purpose the relevant feature is the role of joint acceptance. I work with the following analysis of reasons according to JAA.

*A group has reason r iff all operative members are willing to take a reason r as a group reason to believe p.*⁶

I specified the relevant members as *operative members* following Tuomela (2013, pp. 86-87). Operative members are those that play a role in a group's decision-making, whereas non-operative members do not impact the group's decision-making. This is an important qualification, because it seems implausible that in all cases every member has to be part of the joint acceptance procedure. Some members will not be involved in decisions, but groups

⁴ I remain rather neutral on the notion of 'reason' and 'evidence'. However, given JAA some incompatible notions are ruled out. Reasons cannot be factive, for instance.

⁵ Or would properly express openly (Schmitt, 1994, p. 265).

⁶ The similar analysis in terms of evidence: A group has evidence e for proposition p iff all operative members are willing to accept e as group evidence for p.

with such members can still have reasons that justify beliefs. Hence, reasons under JAA are only determined by operative members.⁷

With the central feature of JAA on hand I can look at Lackey's argument against JAA. Her main worry is that acceptance is an attitude that can be formed at will, and that forming acceptances in an epistemically arbitrary way will lead to justification at will if we accept JAA. After all, group reasons are determined by joint acceptance. All it takes for a group to have a reason *r* that supports a belief that *p* is that the operative members are willing to take the reason *r* as a group reason. Group members can do so for any reason they deem appropriate. It is their choice to take a reason as a group reason. Lackey makes the problem apparent with two cases:

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

⁷ As Schwengerer (2021, p. fn3) hints at, the distinction between operative and non-operative members leads to interesting cases if we do not demand a condition that members express their willingness to jointly accept a reason openly. This would allow for cases in which the operative members determine the reasons for the group, but non-operative members would not be aware of these reasons at all. I will bracket this issue.

smoking and lung cancer and heart disease are unreliable. (Lackey, 2016, p. 351; 2021, pp. 64-65)

In both cases the group members are fully aware of the dangers of smoking. But in **IGNORING EVIDENCE** the group members decide because of financial reasons that they are not going to take the dangers of smoking as a reason for the group to put warning labels on their product. And in **FABRICATING EVIDENCE** they even go a step further and decide for financial reasons to jointly accept that the scientists are liars. Hence in **FABRICATING EVIDENCE** the group has a reason not to believe the results of the scientific studies about the dangers of smoking.

Both cases are in line with JAA. If joint acceptance is all that is required for group reasons, and joint acceptance can be achieved arbitrarily by the group members on any basis, then Lackey's examples seem unavoidable. Groups can generate justification at will. And justification at will also means that we give up on justification being a guide to truth. This poses the *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, p. 353; 2021, p. 67)

Severing epistemic justification and truth-conduciveness is in itself enough of a motivation to avoid IMEP. However, there are also practical concerns in the background. Lackey is worried that if group justification is too easy to achieve we will be unable to blame groups for their beliefs. In both examples Philip Morris turns out to be justified in believing that no labels are necessary, if we follow JAA. But intuitively Philip Morris should be blamed for their beliefs in both **IGNORING EVIDENCE** and **FABRICATING EVIDENCE**. It seems natural to blame them for having beliefs that are epistemically improper – that are unjustified. An account of group justification ought to make good on these intuitions. Lackey proposes that JAA cannot do that. However, in earlier work (Schwengerer, 2021) I disagree and suggest that JAA can give the right verdict. In the next section I present my 2021 response before showing that it comes with consequences that are at least very close to giving up on JAA.

3. The Containment Strategy

Lackey's main concern is that JAA allows for illegitimate manipulation of evidence that leads to severing the connection between group justification and truth-conduciveness. Schwengerer (2021) agrees with that concern, but aims to show that we can rule out illegitimate manipulation of evidence for groups without giving up on JAA. The idea is to accept JAA only as an account of prima facie justification and to use the concept of epistemic expectations taken from Sanford Goldberg (2018) to find a path from prima facie to ultima facie justification. The relevant epistemic expectations are *evidential expectations* – expectations on the body of evidence that an epistemic agent has. They determine what evidence a group should have in a particular situation and thereby give rise to the phenomenon of normative defeat.⁸ If the group's actual evidence diverges from the evidence the group should have the group's ultima facie justification is appropriately reduced. The evidence the group should have provides an upper limit to the group's ultima facie justification. Hence, the epistemic expectations we have towards a group ensure that illegitimate manipulation of evidence is defeated and groups cannot illegitimately manipulate ultima facie justification.

Let me illustrate the idea with IGNORING EVIDENCE: In IGNORING EVIDENCE the group's prima facie justification is determined by the group's reasons. These reasons are established by joint acceptance of the group members. None of the members is willing to take the scientific studies about the danger of smoking as a reason for the group. The group's body of evidence does not include that smoking is dangerous. And neither includes the body of evidence that this danger is a reason to believe that warning labels are necessary. However, in Schwengerer (2021) I argue that it is clear that Philip Morris *should* include both of these propositions in their body of evidence. Philip Morris should know about the dangers of smoking. If Philip Morris believes that warning labels are not required then such a belief is normatively defeated because Philip Morris should know about these studies. We expect Philip Morris to know these studies, because the individual members of Philip Morris know about the studies. We know

⁸ Normative defeat is independently accepted by Lackey (2005; 2016), hence she cannot reject the proposal by rejecting normative defeat. Other proponents of normative defeat include Harman (1980), Kornblith (1983), DeRose (2000), Meeker (2004) and Goldberg (2016).

that the evidence is in principle available, but the group members decided not to take it up as reasons for the group.

IGNORING EVIDENCE functions similar to Kornblith's (1983) well-known physicist case for normative defeat. Imagine a physicist who tunes out his colleagues' comments as soon as they come close to providing him counterevidence to his pet theory. He never goes to talks that might include counterevidence, and neither does he read or engage with scientific articles that have a chance of not supporting his pet theory. In this case the physicist's own body of evidence will end up supporting his pet theory to a high degree. But only because it has been curated in an illegitimate way. There is evidence the physicist should have. He should have listened to colleagues and read relevant papers even if they might not be in favour of his theory. His current belief is epistemically improper. It is unjustified. But not because his body of evidence does not support it, but because of the evidence he should have. His belief is normatively defeated by this evidence he should have.

My solution to IGNORING EVIDENCE in Schwengerer (2021) is supposed to work the same way. Philip Morris' belief is normatively defeated by the evidence the group should have. Moreover, a generalised form of the proposal should also solve FABRICATING EVIDENCE in favour of JAA. Here the idea is not that there is evidence missing that the group should have, but rather that the group has evidence it should not have.

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. (Schwengerer, 2021, p. 12)

This is an explicit step away from Goldberg's (2018) work on normative defeat and epistemic expectations. Goldberg is only interested in cases in which an agent lacks evidence the agent should have, but he does not discuss cases in which an agent has evidence it should not have. Schwengerer (2021) includes both in an attempt to deal with Lackey's counterexamples to JAA. I proposed the following *Evidential Ceiling Principle*:

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). (Schwengerer, 2021, p. 13)

This principle provides the general answer to cases such as IGNORING EVIDENCE and FABRICATING EVIDENCE. The group justification is determined by looking at both the group's evidence and the evidence the group should have. The group's evidence is based on joint acceptance as proposed by JAA and thereby under voluntary control. This is the group's prima facie justification. But the evidence the group should have is determined by our evidential expectations towards the group and thereby not under voluntary control. And this evidential expectation sets the ceiling for the group's ultima facie justification. **ECP*** rules out illegitimate manipulation of ultima facie justification. It ensures that the ultima facie justification is legitimate. It contains the epistemic badness that might be part of the group's prima facie justification – hence I call this proposal the *containment strategy* (Schwengerer, 2021, p. 6).

Clearly the concept of an epistemic expectation performs the heavy lifting in the containment strategy. For the strategy to work we need to make proper sense of the evidence that a group should have at time t. But in this regard Schwengerer (2021) does not give us much to work with. While I make plausible that the norms related to epistemic expectations have to be the same for groups and individuals, this does not provide much help in determining epistemic expectations for groups. The closest I come in my early attempt is in the following two passages:

We have general expectations for individuals that include having acquired adequate evidence for a 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. (Schwengerer, 2021, p. 18)

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. (Schwengerer, 2021, p. 18)

It is not clear at all how exactly we should imagine such a single individual in the position of the group. At best this proposal is underdeveloped, but it might be worse and a step in the completely wrong direction. Granting that my suggested containment strategy works against Lackey's examples in principle, I am now going to fill in the gap by suggesting a way of determining the evidence that a group should have. My approach will diverge from imagining a single person in place of the group. Instead I suggest that Goldberg's (2018) discussion of epistemic expectations gives us a better path to a group's epistemic expectations by looking at the group's social position and the members' beliefs. However, this path will lead us further away from JAA and closer to Lackey's own account of group justification. Perhaps too close for proponents of JAA and too close to fulfil the aims laid out in Schwengerer (2021).

4. Containment and Evidential Expectations

Let me start working towards an account of evidential expectations for groups by looking at Goldberg's (2018) explanations of epistemic expectations. These expectations come in two forms: basic epistemic expectations and non-basic (or general) epistemic expectations. Basic epistemic expectations determine what Goldberg calls the core criteria for epistemic propriety. They provide us with a notion of prima facie justification. Goldberg understands these core criteria mostly in terms of process reliabilism. An agent satisfies the core criteria only if their belief-formation process is truth-conducive (Goldberg, 2018, p. 151). Moreover, Goldberg also includes a condition that demands an agent to consider information already in its cognitive system. This is spelt out as a coherence and updating condition for an agent.

[...] we also have expectations regarding how others adjust their beliefs in the course of acquiring new beliefs— in other words, how they adjust to newly acquired evidence. Thus we expect people to adjust their background beliefs to newly acquired evidence in a way that reflects the twin interests in acquiring truths and avoiding errors. [...] We expect them to be sensitive to incoherence in their own belief system, and to adjust their beliefs when they do detect incoherence. (Goldberg, 2018, p. 152)

The agent's beliefs have to cohere with the information available to the agent – the information available in the agent's cognitive system. If new information comes in the beliefs have to be adjusted to satisfy coherence. The upshot here is that we have expectations that

agents use reliable belief-forming processes and that agents have beliefs that cohere with the information available to their cognitive systems. These are the basic epistemic expectations.

When an agent meets the basic epistemic expectations for a belief that belief is *prima facie* justified. This might already seem problematic for groups and JAA, given that JAA itself does not provide any form of reliability for the group. However, in my previous work (2021, pp. 14-17) I offer a solution borrowing from Schmitt (1994) by pointing to a group's *charter* as its code of conduct and its *office* as its guiding aim.⁹ This restricts the joint acceptance done by the group members. Each group has a process specific to the group that leads to joint acceptance. This process is captured in the charter. In a simple case the charter captures how deliberation and discussion in a group leads towards joint acceptance. It determines what kind of conduct in deliberation is permissible, and what is not. For instance, a charter might make violence as a path towards joint acceptance inadmissible. The charter itself determines the group's office and structures how the group aims to achieve that office. Office and charter need not be written down or be made explicit at all. But the group members need to be aware of both office and charter, such that their individual actions can be guided by the office. I do not provide a story how office and charter come about. However, given that JAA is a close relative of joint commitment accounts of collective intentions (Gilbert, 1989; 1990; Tuomela, 1991; 2005) it seems natural to assume that charter and office are established by some form of joint commitment of operative members for proponents of JAA.

In Schwengerer (2021, pp. 14-17) I suggest that we can individuate a group's belief-forming processes in terms of the charter and use that to determine their overall reliability. If a group following a specific charter gets things right often enough, the group satisfies the basic epistemic expectations. The charter can satisfy reliability even if it allows for a somewhat epistemically arbitrary process of forming reasons and beliefs in individual cases. Reliability is individuated on a very broad level of the group forming beliefs generally and therefore only gives minimal constraints on JAA that satisfy the core criteria. A group's charter can be such that it is overall reliable, even though it includes somewhat epistemically arbitrary belief-formations in particular subject areas. For instance, Phillip Morris can satisfy the core criteria overall even if they base their beliefs about the dangers of smoking on economic reasons (Schwengerer, 2021, p. 17). What we end up with is a first amendment to JAA: *a group's*

⁹ See also Kallestrup (2020).

reasons only provide prima facie justification if the belief formation by the group following a charter and office is overall reliable.

This is not the only way to go. Another option would be to use different constraints for Goldberg's core criteria and use a more internalist variant for core criteria. For instance, we might merely opt for proportioning one's beliefs to one's evidence as the core criterion, which might be more palatable than individuating the group's belief-forming process as widely as suggested. To save space I go with the reliabilist version and point to the aforementioned discussion in Schwengerer (2021).

With the basic epistemic expectations satisfied I can move on to the containment strategy. The strategy makes sure manipulated evidence that provides prima facie justification does not constitute ultima facie justification. For this we need to look at the second type of epistemic expectations proposed by Goldberg: non-basic (or general) epistemic expectations. These determine whether a prima facie justified belief is also ultima facie justified. Whether these non-basic expectations are satisfied determines whether the belief is all things considered epistemically proper or not. And non-basic expectations are closely related to a social position and role of an agent. What we expect epistemically from other people is not universally the same for everyone in every circumstance. Consider a situation in which a doctor informs us of potential side-effects of a medical procedure. We expect that the doctor is well-informed about relevant studies about the procedure and its side-effects. If there are numerous studies highlighting a particular risk with that procedure we expect them to know these studies. It's part of their job and responsibility as a doctor! However, if I was just talking to a carpenter friend about the same medical procedure and potential side-effects, then there would not be any expectation about them being acquainted with scientific studies. If the doctor misses a relevant study I will blame them. If my friend is completely unaware of a relevant study blaming them would be ludicrous.

Already in this example it is clear that in assessing the doctor's belief we do not only care about the basic epistemic expectations being met. The doctor can be reliable and their beliefs coherently updated to the information in their cognitive system, but nevertheless the doctor's beliefs about the side-effects would ultimately be epistemically improper. They are improper because the doctor fails to satisfy the non-basic epistemic expectations. Given their social role we expect them to have looked at relevant studies, but they have not. They lack a particular

piece of evidence that they should have. Hence Goldberg claims that “It is when she fails to have e.g. the evidence she should have had (given the role(s) she is playing in her epistemic community/ies), that considerations of social epistemic responsibility can disqualify an otherwise epistemically proper belief of hers from so counting” (Goldberg, 2018, p. 188). This is the idea of a non-basic epistemic expectation in a nutshell. It is an expectation about the body of evidence an agent ought to have in a particular situation given their particular social role in a community.

This brings me a step closer to being able to determine what evidential expectations are present for a group such as Philip Morris in *IGNORING EVIDENCE* and *FABRICATING EVIDENCE*. The question can be answered by looking at the particular social role that a group like Philip Morris has. The group’s social role determines what sort of evidence the group is expected to have, and the individual group members determine what particular pieces of evidence that are of this sort are easily available and therefore should be in possession of the group. Philip Morris’s role as a tobacco company comes with expectations of having evidence about the safety of their tobacco product in general, and the group members’ knowledge of the studies about the dangers of smoking gives us the particular evidence that Philip Morris should have because the evidence is both easily available and relevant to Philip Morris’s social role.

My aim now is to develop this picture in more detail. Let me start with the social role of a group. A group’s social role is partially determined by the group and partially by the community the group is in. The former stems from a group setting its own agenda in the context of a community. I will discuss limits set by the community later, for now let me focus on the group’s agency in determining their social role. Groups – as opposed to mere collections of people¹⁰ – form with some goal or aim the group members want to achieve collectively. Even short-lived temporary groups commonly have an aim. For instance, if three people work together to push a broken car off the street they will form a group with the purpose of getting the car out of traffic. And the members each do their part to achieve this aim.

I have already introduced the tools required to capture this phenomenon properly: a group’s charter and office – its code of conduct and guiding aims. Charter and office constitute group

¹⁰ This distinction is similar to Ritchie’s (2015) distinction between Type 1 (structured) groups and Type 2 (feature) groups.

internal factors that impact the non-basic epistemic expectations we have towards the group. The group's goal and the group's means to achieve this goal determine what kind of evidence the group ought to have – namely evidence that is relevant for the goal and the means of achieving the goal. Take Philip Morris with an office that includes producing and selling tobacco product to earn a profit. This office gives us an idea of which kind of evidence is relevant for Philip Morris. Or rather, it gives us a clear path to tell which kind of evidence must be relevant for Philip Morris. A piece of evidence must be relevant if it includes information that impacts the production of tobacco products and the sell of these products. After all, that is what the group aims for and tries to achieve. Insofar as the group intends to achieve its office practical rationality requires the group to be interested in information relevant for its office. This is one source for evidential expectations towards the group. The group locates its own role in the community via its charter and office and in doing so generates legitimate epistemic expectations. Philip Morris identifies itself as a tobacco company and thereby generates the epistemic expectations that Philip Morris takes in the available evidence about tobacco. By claiming a domain in which Philip Morris will be active to achieve its office Philip Morris also claims sufficient expertise in that domain to achieve its office and takes on the responsibility of said claim.

Charter and office are a group's attempt to determine its own social role in front of the background of a community. But a group cannot arbitrarily pick any social role with its epistemic expectations. A group can only pick a social role that is available within the framework a community provides. In the case of Philip Morris the community offers the option of a particular social role as a company producing, buying and selling products, and Philip Morris can take up the offer. For companies this can become formalised in various forms of registers, such as a company or trade register. Similarly clubs, societies or unions can be formalised to make their social role explicit. However, even informal groups take up a social role in a community that is partially determined by the group's choice and partially determined by the social roles offered in a community. A group of amateur football players that regularly meets at a public football field does not have an explicit charter and office, but nevertheless the group takes up a social role in a community. For instance, the community can legitimately expect that the group knows when the weather conditions are such that playing football would damage the field. And the community can also legitimately expect that the group's evidence will include the current weather conditions. These expectations come

with the social role of being an informal football team on a public football pitch. Here the group chose its charter and office without making them explicit. In doing so they have to take up a social role within the larger community that is not completely up to the group.

These community factors determine non-basic epistemic expectations from outside the group. As a community we have set particular rules and expectations for specific social roles. Especially for those social roles that can be potentially harmful to people and objects in the community. Take again Philip Morris with an office that includes producing and selling tobacco products. If it was only up to Philip Morris what evidence counts as relevant for their office, then the long-term health impacts of the product would likely not show up as relevant. As long as their customers are not dying out and impact their bottom line Philip Morris might not care. However, the larger community clearly cares about health impacts of any product offered on the marketplace. We hold companies responsible for any such health risks constituted by their product. Part of the social role of a company that sells goods is that the seller has to guarantee that a product is safe to use. Plausibly this has been part of the social role of vendors and retailers for a long time, but has become more and more formalised. Nowadays we find among other regulations health codes for food and extensive testing requirements for new pharmaceutical products. Moreover, consumers have a claim right on knowing the ingredients of products in order to make judgment on their safety.¹¹ These rights, demands and regulations come with the social roles that vendors, retailers and merchants have, regardless whether they are individuals or groups. If this is correct, then the social role that requires a seller to guarantee that a product is safe also provides a source for epistemic expectations. For if a seller guarantees that a product is safe to use, the seller also gives rise to epistemic expectations about the seller's evidential basis for that guarantee. The seller has to be epistemically responsible in order to give a guarantee that satisfies the demands of their social role. The seller has to take available, relevant evidence into account. In case of Philip Morris this entails that we are entitled to expect Philip Morris to have relevant evidence about the safety of their tobacco product. Philip Morris should know about available studies showing that smoking is dangerous.

I have now shown why we have non-basic epistemic expectation towards groups. These expectations are about the kind of evidence we can legitimately expect a group to have, if

¹¹ For the idea of epistemic claim rights see Watson (2018).

evidence of that kind is available. For instance, we expect a tobacco company to have relevant evidence about tobacco and its impact on our health, if such evidence is available. What I still require is an explanation of when exactly a piece of evidence is available for a group. Only then can I explain what piece of evidence a group is expected to have. And I need that explanation in order for any version of my containment strategy to work.

The idea of normative defeat is closely related to the thought that some evidence is easily or readily available for an agent. Kornblith's (1983) physicist has access to evidence from colleagues, could read the relevant articles and could attend conference talks. In some cases, the physicist has to put in a remarkable effort in order to avoid the evidence. In other cases, it would be at least very easy to access the evidence. Perhaps as easy as opening a page in a book that is already in front of him. For a piece of evidence to constitute a normative defeater the evidence has to be easily available, as in the example. It is difficult to provide a universal delineation of instances in which evidence is easily enough available, and instances in which it is not. However, all I need is to show that there are clear cases of evidence that is sufficiently easily available for groups when the group members individually have such evidence.

Let me start by illustrating the idea with IGNORING EVIDENCE. In IGNORING EVIDENCE each of the operative members is individually aware of the massive amounts of scientific evidence linking smoking to cancer. And each individual member believes that is a reason to believe that warning labels should be placed on the cigarette boxes. So each individual member has evidence supporting the belief that warning labels are called for. My suggestion is that this evidence the individual members have is easily available for the group. The group's reasons are determined by joint acceptance of the operative group members. But with every operative group member being in possession of the member evidence m , and playing a part in determining the group's evidence g , it seems obvious that the member evidence is also available for joint acceptance by the group. Moreover, it is not just merely available – because everything is in some sense available for the group to accept as a reason. It is available as a default option. Group members need a reason not to jointly accept m as a group reason, when every operative member includes m in their evidence. This can be established in two ways. First, the charter of any successful group requires generally truth-conducive processes of forming group reasons. For a group to successfully fulfil their office it cannot be wholly unreliable. Otherwise the aims of the group cannot be reliably achieved. The actions of the

group are guided by group reasons and if these group reasons are not sufficiently hooked up to the world the actions will end up in frequent failure. Building on member evidence as a default for joint acceptance is a straightforward way to achieve truth-conduciveness if the group members are reliable. Second, in the dialectic of the containment strategy I am already presupposing for my discussion that the basic epistemic expectations are satisfied, which I addressed earlier pointing to Schwengerer (2021). Hence, I restricted JAA to groups with a charter and office that leads to overall reliable belief-formation. The group's general belief-formation method is already set as reliable and this can only be the case if the default option for forming group reasons is truth-conducive.¹² The best explanation for this truth-conduciveness is that group members are willing to jointly accept a reason as a group reason by default if their own evidence includes that reason. In this way the truth-conduciveness of belief-formations of individual operative members will transfer to the group.

This already might look suspiciously close to a form of reason-aggregation and far removed from the idea of pure joint acceptance for group justification. However, it is still a JAA as long as diverging from the default still leads to prima facie group justification. All the default option entails is that evidence held by individual members is easily available for the group. Not taking it up as evidence for the group requires more effort than accepting it – just like the physicist had to put in additional effort to look away from evidence.

The picture gets slightly more complicated when there is no uniform evidence basis shared by operative members. In IGNORING EVIDENCE every member had evidence about the danger of smoking. But often members diverge in their evidence. Some have a particular piece of evidence, others do not. Perhaps different members even have conflicting, incompatible evidence. Clearly, if they have conflicting evidence that persists even after full disclosure and deliberation without any majority that evidence cannot constitute a default joint acceptance option.¹³ Hence, this evidence will not be available for the group. I suggest the following first pass:

¹² This does not entail that every instance of a group jointly accepting a reason to believe a proposition is truth-conducive, just that overall their process is. Schwengerer (2021) already makes this assumption calling it a “flat-footed” reliabilist assumption (p. 15).

¹³ Cf. Lackey (2021, chapter 2.5 and 2.6) for a thorough discussion of problems that can occur with conflicting bases.

Easy Evidence – First Pass: Evidence E is easily available for the group if and only if a significant percentage of the operative members of G (a) have evidence E, and (b) are such that adding together their evidence yields an evidence set that is coherent.¹⁴

However, this first pass is not good enough. It captures sufficient conditions for evidence that is easily available for a group, but not necessary conditions. If enough group members have evidence, that evidence is easily available for the group. But in some cases evidence is easily available even if no group member has a particular piece of evidence. This happens if some piece of evidence is itself easily available for enough group members. Remember Kornblith's physicist who tunes out any colleagues who might propose counterevidence. Now suppose a group consists of only members akin to this physicist. The group has only members who lack evidence that they should have. Take the following example:

DEFEATED MEMBERS: Philip Morris is one of the largest tobacco companies in the world. Each of its operative members is presented with a variety of newspapers about tobacco products each morning. However, each operative member only opens those sources that never feature any information about dangers of smoking. Some of the other newspapers regularly feature scientific studies about the dangers of smoking, but the operative members never look at these. Suppose that each member also avoids all other sources of information that might include articles about negative health effects of smoking. Hence, each individual member believes that smoking is not dangerous at all. All these members would properly express a willingness to accept that Philip Morris's position is that smoking is perfectly healthy.

In **DEFEATED MEMBERS** every single operative member has a normative defeater and is therefore not justified to believe that smoking is healthy. They have a normative defeater because they have easy access to information about the dangers of smoking, and their social role comes with evidential expectations related to the dangers of smoking. We expect operative members of a tobacco company to be aware of easily available evidence about the safety of their product. If we just take the first pass for Easy Evidence from above Philip Morris would not have easily available evidence that smoking is dangerous in **DEFEATED MEMBERS**. After all, no operative member has evidence that smoking is dangerous. Hence, Philip Morris

¹⁴ **Easy Evidence – First Pass** is inspired by one of Lackey's (2016; 2021) conditions for group justification. This connection will be further highlighted in section 6.

would not have a normative defeater for their belief that smoking is safe. But that seems to be the wrong verdict. The right verdict in DEFEATED MEMBERS is that both the operative members and Philip Morris lack ultima facie justification. The evidence is easily available to Philip Morris because it is easily available to the operative members of Philip Morris. Hence, **Easy Evidence** requires a further adjustment:

Easy Evidence – Second Pass: Evidence E is easily available for the group if and only if a significant percentage of the operative members of G either (a) have evidence E, and are such that adding together their evidence yields an evidence set that is coherent; or (b) should have evidence E¹⁵ and are such that adding together the evidence they should have yields an evidence set that is coherent.

As a reviewer pointed out, this is still not sufficient for a general version of **Easy Evidence**. This second pass works for groups in which all operative members are equal. Their word has the same weight, and they all are operative for any decision or belief formation. But many groups are more structured than that. For instance, suppose the finance officer of a company has information about the company's finances, but the other five executives do not have that information. It seems plausible that the group should have that information as evidence, even though the other five operative members do not have that information. Hence, in such a case a low percentage of operative members have the evidence individually, but the group nevertheless should have it.

To deal with this issue I have to introduce more moving parts to **Easy Evidence – Second Pass**. The first way to go is to relate operative members to subject matters or domains. Who is an operative member of a group can change depending on what the group is dealing with at the moment. Someone in charge of human resources will be an operative member in hiring decisions, but not in decisions about production issues. This can deal with some instances in which only few members having evidence individually can be sufficient for something to be easily available evidence for the group. In some domains these few members by themselves constitute a significant percentage of the operative members. However, this will not be enough for the case of the finance officer, because in plausible cases the other executives will always have some say, even in financial matters. They will be operative members, regardless

¹⁵ Where should have evidence E for individuals is understood in terms of evidential expectations towards an individual in their role as a member of the particular group. See Goldberg (2016; 2018).

whether the finance officer is the main person in charge of finances. To deal with this problem one has to look at the operative group members as having different roles in the group's decision-making. Not every operative member counts the same with regard to every subject matter or domain. This can be captured by giving the operative members different weights in the decision-making process. Financial issues fall clearly under the domain assigned to the finance officer. Hence, their actions and decisions have more weight than the decisions of other operative members.¹⁶ So, instead of merely looking at a significant percentage of operative members to determine easily available evidence, I need to look at a significant percentage of members relative to their weight in regard to the subject matter of the evidence. A low percentage of members with high weight in regard to the subject matter of the evidence can be enough for evidence to be easily available. On the other hand, also a higher percentage of members with lower weight in regard to the subject matter of the evidence can be sufficient. It is easy to see that these considerations become rather messy and heavily depend on the structure of the group.¹⁷ This brings me to the final version of **Easy Evidence**:

Easy Evidence: Evidence E is easily available for the group if and only if a significant percentage of the operative members of G relative to their weight in regard to the subject matter of the evidence either (a) have evidence E, and are such that adding together their evidence yields an evidence set that is coherent; or (b) should have

¹⁶ Having members with different weight might look to be in conflict with joint acceptance accounts. However, following Gilbert these cases can be understood as involving derived, non-basic joint commitments. These non-basic joint commitments depend on a basic joint commitment authorizing a given person or body to create new joint commitments for the whole group in a particular way. For instance, a finance officer can be authorized by a joint commitment to be in charge of finances. In such a case the weight of the finance officer's opinion with regard to financial issues is determined by a prior joint commitment by the operative members.

¹⁷ One attempt to make this clearer would be to model it more formally by assigning every operative member a weight between 0 and 1 for a subject matter or domain such that the combined weight for all operative members is 1. Suppose the group has 6 operative members. One member is the finance officer who has the information and has the most weight in financial decisions (0.75). One other operative member also has the information, but has little weight (0.05). The other 4 operative members do not have the information and also have little weight in financial decisions (0.05). We then add the weights of the operative members that have (or should have) the evidence. In this case this is simply the sum of one member with weight 0.75 and one member with 0.05 for a total of 0.80. The closer to 1 this calculation gets, the easier the evidence is available for the group. It reaches 1 when every operative member has or should have the evidence individually. It reaches 0 when no operative member has or should have the evidence individually.

evidence E^{18} and are such that adding together the evidence they should have yields an evidence set that is coherent.

Two remarks on **Easy Evidence**: First, it is important that I am restricting this to evidence that is easily available. Plausibly much more evidence is available¹⁹, but in order to pinpoint the sources of evidential expectations and normative defeat I am only picking out evidence that is without a doubt *easily* available. And, second, the evidence individual members should have also relates to the ability to fully disclose and deliberate between the group members. Other group members are also easily available sources of information that determine what evidence a member should have. Hence, a group cannot restrict what evidence is easily available to the group by limiting discussion between members.

I can now combine the ideas discussed on social roles and non-basic epistemic expectations and this account of available evidence for a group. The non-basic epistemic expectations are determined as follows:

Expectations: A group G is legitimately expected to have evidence E by a community C iff

- (i) E is relevant for the group in its social role in the community; and
- (ii) E is easily available for the group, where E is easily available for the group if and only if a significant percentage of the operative members of G relative to their weight in regard to the subject matter of the evidence either (a) have evidence E , and are such that adding together their evidence yields an evidence set that is coherent; or (b) should have evidence E and are such that adding together the evidence they should have yields an evidence set that is coherent.

With this account of group expectations in hand I can provide a new account of group justification following JAA and the containment strategy.

5. The Anchored Joint Acceptance Account of Group Justification

Let me start with the core idea of JAA: A group has reason r iff all operative members are willing to take a reason r as a group reason. As a full account of group justification this is not

¹⁸ Where should have evidence E for individuals is understood in terms of evidential expectations towards an individual in their role as a member of the particular group. See Goldberg (2016; 2018).

¹⁹ There is a sense in which any kind of evidence is available with JAA.

functional, because of Lackey's worries about illegitimately manipulated evidence shown in IGNORING EVIDENCE and FABRICATING EVIDENCE. However, with the suggestion of the containment strategy we can hold on to an amended JAA as determining prima facie justification:

Prima Facie Justification: A group G is prima facie justified to believe that p iff

- (i) all operative members are willing to take a reason r as a group reason for p;
and
- (ii) the group's belief-forming processes following its charter and office are overall reliable.²⁰

(i) is the original JAA idea, whereas (ii) restricts it to groups that are overall reliable. It is still a JAA insofar as groups that are overall reliable can generate reasons for p at will, as long as their overall truth-ratio is sufficiently high. This allows Philip Morris to satisfy **Prima Facie Justification** in IGNORING EVIDENCE and FABRICATING EVIDENCE.

While this might be right, it is rarely useful. We are primarily interested in ultima facie justification, not merely the prima facie kind. And here the containment strategy kicks in. While prima facie justification for a proposition can be wholly epistemically arbitrary, ultima facie justification must not be. An account of ultima facie justification has to provide us with the means to distinguish legitimately manipulated evidence from illegitimately manipulated evidence in order to avoid Lackey's argument. The containment strategy does this with an epistemic ceiling principle.

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). (Schwengerer, 2021, p. 13)

ECP* restricts the maximal justification a group can have for a belief that p at time t by pointing to the evidence that the group should have had at t. And while Schwengerer (2021) did not provide us a story of how this evidence the group should have is determined, I have filled this

²⁰ Note that this is a definition for propositional justification, but nevertheless involves a condition about belief-forming processes.

gap in the previous section. The group should have the evidence that is both relevant for the group in its social role and easily available. Easy availability is understood in terms of the evidence that the individual group members have or should have. A piece of evidence is easily available for the group if and only if a significant percentage of the operative members of G relative to their weight in regard to the subject matter of the evidence (a) have evidence E, and are such that adding together their evidence yields an evidence set that is coherent; or (b) should have evidence E and are such that adding together the evidence they should have yields an evidence set that is coherent. This gives us the tools to build an account of ultima facie group justification. For this we have to adapt the account for prima facie justification to include degrees of justification and then combine it with the ceiling principle.

Prima Facie Justification to Degree d' : A group G is prima facie justified to believe that p to a degree d' iff

- (i) all operative members of G are willing to take evidence E as group evidence to degree d' for p; and
- (ii) the group's belief-forming processes following its charter and office are overall reliable.

Ultima Facie Justification to Degree d : A group G is ultima facie justified to believe that p to a degree d iff

- (1) there is some degree d' , such that G is prima facie justified to believe that p to the degree d' ; and
- (2) there is some degree d^* , such that d^* is the degree of justification for p the group G should have according to G's social role and the evidence easily available to G (and updated in a rational fashion); and
- (3) d is the minimum value of $\{d', d^*\}$.²¹ If d' equals d^* both are minima and d equals both d' and d^* .

Because d is determined by picking the minimum value and both values considered are the same whenever G's actual evidence matches the evidence G should have, the degree of ultima

²¹ (1)-(3) together are a reformulation of the epistemic ceiling principle ECP*.

facie justification d equals the degree of prima facie justification d' , whenever G 's actual evidence matches the evidence G should have (that is: $d=d'=d^*$).

Condition (1) starts with a degree of prima facie justification. Condition (2) determines the degree of justification the group should have based on the non-basic epistemic expectations. Finally, (3) then determines the degree of ultima facie justification. Whenever the group's actual evidence does not match the evidence the group should have we end up with a set containing two elements with different values. (3) then picks the minimum value – in this case the lower value of the two elements of the set. However, when the evidence the group has and should have match, then both elements have the same value and both determine the minimum value. Hence, in this case (3) does not lower the degree of justification by going from prima facie to ultima facie justification. The degree of ultima facie justification d then equals the degree of prima facie justification d' .

To make this easier to grasp consider two examples with concrete values. Suppose G 's degree of prima facie justification d' is 0.8 based on G 's actual evidence. Furthermore, suppose that the degree of justification based on the evidence G should have was 0.6. (3) then looks for the minimum value of the set $\{0.8;0.6\}$. In this case the minimum value is 0.6. This determines the degree of ultima facie justification d as 0.6.

In a case in which G 's actual evidence matches the evidence G should have, (3) gives us a set with two elements with matching values that are both minima – for instance $\{0.8;0.8\}$ – and d is determined by the value of any of those elements. In this case d would be 0.8, which matches the degree of prima facie justification d' .

We can also add a condition for outright justification in case we want to work with a binary notion of justification.²² A group G is outright justified to believe that p iff (1)-(3) holds, plus an additional fourth condition:

(4) The degree of justification d surpasses a threshold for outright justification.

Condition (4) looks at whether the resulting ultima facie justification from (1)-(3) meets a threshold for outright justification. (4) is only required if we want to make binary judgments

²² I use 'outright justification' to capture a binary talk of justification. We sometimes speak of beliefs as being either justified or not justified, rather than in terms of degree of justification.

about justification. Otherwise we can also simply end with (1)-(3) that give us a degree of ultima facie justification for the group belief.

This is my proposal for ultima facie group justification – *an anchored joint acceptance account of group justification*. It is a form of JAA that is normatively anchored in the non-basic expectations generated by the group members. It is anchored because mainly the evidential states of group members determine which evidence the group should have. But it is still a version of JAA because these expectations can only lower justification. They cannot generate justification. The only way to generate justification is the group members jointly accepting a piece of evidence as evidence for the group.

The account gives us the correct verdict for IGNORING EVIDENCE and FABRICATING EVIDENCE. In both cases the group has prima facie justification. But in neither case the group's evidence is identical with what it should be. Hence condition (3) kicks in and decreases the justification. In both cases the ultima facie justification will be too low to pass any reasonable threshold for outright justification. Moreover, my account still allows for the same explanation of divergence between group justification and group member justification that has been suggested by Schmitt (1994) and Hakli (2011). The account can capture how, for instance, a jury's justification can come apart from the justification individual jurors have. The group can even be justified if no individual member is. Hence, the account keeps the upside of traditional joint acceptance accounts while getting rid of their main problem.

One might be still worried that not all of Lackey's criticism has been fully addressed. According to my account Philip Morris does not have any reason to believe it should put warnings on cigarette boxes in IGNORING EVIDENCE. Lackey on the other hand can explain how Philip Morris has such a reason, but fails to act on that reason in an appropriate way. One might take that to be an advantage of Lackey's view over mine. However, I do not think that is an effective objection to my proposal, because my account still explains that Philip Morris is doing something epistemically wrong here – something that has severe consequences. My account still gives us the result that we can and should hold Philip Morris responsible for the ill effects caused by smoking. This is explained by the verdict that Philip Morris should have these reasons. Not having these reasons is a failure of Philip Morris. And both this failure and all its consequences are the responsibility of Philip Morris.

Lackey and I agree that something is going wrong in the Philip Morris case. And both Lackey and I agree that Philip Morris can and should be held responsible. The only difference is where we locate Philip Morris's failure. Lackey thinks that in order to hold Philip Morris responsible we have to ascribe a particular reason to Philip Morris that they do not act on appropriately. I suggest that is not the case. We can hold Philip Morris responsible because they *should have a particular reason*, but do not have it. More precisely: Philip Morris should have this particular reason because it is so easily available to Philip Morris. That Philip Morris still does not accept that reason is something that we can hold Philip Morris responsible for. I think that is the right result and it matches our intuition that Philip Morris is blameworthy for doing something epistemically wrong. Hence, I am willing to bite the bullet: Philip Morris does not have a reason to put warnings on cigarette boxes. But it is only a small bullet to bite, because on my account Philip Morris should have such a reason. Not having that reason is blameworthy and prompts the intuitions in IGNORING EVIDENCE.²³

6. Lackey's Argument Revisited

The problem of illegitimately manipulated evidence threatened JAA because JAA's arbitrariness seemed to lead to unpalatable consequences. Following the containment strategy these worries can be answered with my anchored joint acceptance account of group justification. But my proposal might not make proponents of JAA all that happy. My view heavily relies on the role of the individual group members and their evidential states. If the containment strategy was meant to defend JAA, then this might be more of a pyrrhic victory. This will become clear when we look at my results and compare it to Lackey's own proposal for group justification. She calls her view the *group epistemic agent account*:

A group G justifiably believes that p if and only if

²³ Here is also a good point to highlight that in some cases a group should believe p while the group believes not-p, and in some other cases a group should believe p while the group has no belief about p whatsoever. IGNORING EVIDENCE is neutral in this regard, because that difference does not matter for Lackey's criticism of joint acceptance accounts. My account explains the justification of a group belief that p if the group has such a belief that p. However, my broader discussion in this paper also explains whether the group should believe p, even when the group does not hold any belief with regard to p. What the group should believe is determined by the evidence the group should have. And the evidence the group should have is determined by the evidential expectations on the group, based on the group's social role and easily available evidence. Thanks to a reviewer for emphasizing this point.

(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. (Lackey, 2016, p. 381; 2021, p. 97)

In Lackey's account (1) is supposed to be a response to the worry of manipulating evidence in JAA. (1) relates the group's justification to the justification of individual group members and thereby to the world. Group justification becomes truth-conducive. Lackey's (1) is very close to what I suggested in the first pass for **Easy Evidence**. The main difference between Lackey's account and my proposal is that the group member states determine group justification directly in Lackey's account, whereas they only determine non-basic epistemic expectations in mine. But they are meant to do the same job: rule out illegitimate manipulation of evidence by anchoring group justification to member justification. This shows the force of Lackey's argument against JAA. Even if the containment strategy is successful JAA has to accept some form of anchoring condition. Any successful form of JAA has to take some part of Lackey's proposal on board. Some proponents of JAA might not be willing to accept any kind of anchoring to group members, even if it is only anchoring as determining conditions for normative defeat. My strategy in previous work (Schwengerer, 2021) was not meant to lead straight back to an anchoring account. But I believe it is the consequence of spelling out non-basic epistemic expectations properly, and gives us a viable candidate for a view of ultima facie group justification.

I want to end by making the view slightly more palatable for friends of JAA. Lackey claims that her argument against JAA leads to a dilemma for JAA. She writes that "[...] either group reasons are determined by joint acceptance or they are not. If they are, the view succumbs to the Illegitimate Manipulation of Evidence Problem. If they are not, the view is not a joint acceptance account. Either way, inflationary nonsummativism is left wanting" (Lackey, 2016, p. 354). As I have said in the previous section, my account does not fall into either of the horns.

It is still a joint acceptance account because justification is only generated by the group members jointly accepting reasons as reasons for the group. But the view does not succumb to the problem of illegitimate manipulation of evidence, because the evidence the group should have is determined by the group members. Both joint acceptance and available evidence play a role in determining ultima facie justification. Maybe this is still enough for proponents of JAA.

Acknowledgments

Thank you to two anonymous reviewers for their helpful comments and recommendations. An early version of the paper was presented at a research seminar at the University of Duisburg-Essen led by Raphael van Riel. Thank you to all participants for their feedback. Further thanks to Lars Dänzer for detailed comments on a draft version. Work on the article has been funded by Deutsche Forschungsgemeinschaft (Project 462399384). Open Access funding enabled and organized by Projekt DEAL.

References

- Bird, A. (2010). Social knowing: The social sense of 'scientific knowledge'. *Philosophical Perspectives*, 24(1), 23-56.
- DeRose, K. (2000). Ought We to Follow Our Evidence? *Philosophy and Phenomenological Research*, 60(3), pp. 697-706.
- Gilbert, M. (1989). *On Social Facts*. London: Routledge.
- Gilbert, M. (1990). Walking Together: A Paradigmatic Social Phenomeon. *Midwest Studies in Philosophy*, 15(1), 1-14.
- Goldberg, S. (2016). On the epistemic significance of evidence you should have had. *Episteme*, 13(4), pp. 449–470.
- Goldberg, S. (2018). *To the best of our knowledge: social expectations and epistemic normativity*. Oxford: OUP.
- Hakli, R. (2011). On Dialectical Justification of Group Beliefs. In H. B. Schmid, D. Sirtes, & M. Weber (Eds.), *Collective Epistemology* (pp. 119-153). Frankfurt: Ontos.

- Harman, G. (1980). Reasoning and Evidence One Does Not Possess. *Midwest Studies in Philosophy*, 5(1), pp. 163-182.
- Kallestrup, J. (2020). Group virtue epistemology. *Synthese*(197), 5233–5251. doi:10.1007/s11229-016-1225-7
- Kornblith, H. (1983). Justified Belief and Epistemically Responsible Action. *Philosophical Review*, 92(1), pp. 33-48.
- Lackey, J. (2005). Memory as a Generative Epistemic Source. *Philosophy and Phenomenological Research*, 70(3), pp. 636-658.
- Lackey, J. (2016). What Is Justified Group Belief? *Philosophical Review*, 125(3), 341-396.
- Lackey, J. (2021). *The Epistemology of Groups*. Oxford: Oxford University Press.
- Meeker, K. (2004). Justification and the Social Nature of Knowledge. *Philosophy and Phenomenological Research*, 69(1), pp. 156-172.
- Ritchie, K. (2015). The Metaphysics of Social Groups. *Philosophy Compass*, 10(5), pp. 310-321.
- Schmitt, F. F. (1994). The Justification of Group Beliefs. In F. F. Schmitt (Ed.), *Socializing Epistemology: The Social Dimensions of Knowledge* (pp. 257-287). Lanham, MD: Rowman and Littlefield.
- Schwengerer, L. (2021). Defending Joint Acceptance Accounts of Justification. *Episteme*. doi:10.1017/epi.2020.55
- Tollefsen, D. (2009). Wikipedia and the Epistemology of Testimony. *Episteme*, 6(1), pp. 8-24.
- Tuomela, R. (1991). We Will Do It: An Analysis of Group-Intentions. *Philosophy and Phenomenological Research*, 51(2), pp. 249-277.
- Tuomela, R. (2005). We-intentions Revisited. *Philosophical Studies*, 125(3), 327–369.
- Tuomela, R. (2013). *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.
- Watson, L. (2018). Systematic Epistemic Rights Violations in the Media: A Brexit Case Study. *Social Epistemology*, 32(2), pp. 88-102.