

Self-Knowledge in Joint Acceptance Accounts

Lukas Schwengerer (University of Duisburg-Essen)

Penultimate Draft – Please Cite The Published Version:

<https://www.tandfonline.com/doi/full/10.1080/09515089.2024.2346174>

Abstract

This paper closes a gap in joint acceptance accounts of the mental life of groups by presenting a theory of group self-knowledge in the joint acceptance framework. I start out by presenting desiderata for a theory of group self-knowledge. Any such theory has to explain the linguistic practice of group avowals, and how self-knowledge can play a role in practical and moral considerations. I develop an account of group self-knowledge in the joint acceptance framework that can explain these desiderata. I argue that a group has self-knowledge of its own attitude A if and only if (i) the operative members jointly accept that the group has attitude A; and (ii) that joint acceptance is non-deviantly caused by the operative group members' belief that the members jointly accepted attitude A; and (iii) the group has attitude A. Together (i) to (iii) give us conditions for a group having a justified, non-lucky, true belief about its own attitude.

Keywords: Self-Knowledge; Joint Acceptance Account; Group Knowledge; Collective Knowledge

Funding: Work on this article has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 462399384.

1 Introduction

People frequently ascribe mental states to social groups¹. One might say that a reading group intends to read a particular book, or that Microsoft intends to buy Activision Blizzard. These ascriptions are not limited to intention. It is commonplace to ascribe knowledge and belief to groups. Apple knows what their next product is and Google knows a lot about its users. This fact about ordinary linguistic practice has been picked up – among others – by social

¹ Hereafter just 'groups'.

ontologists and epistemologists who provide different attempts to capture what is behind these ascriptions of mental states to groups. In this paper I will only focus on one kind of these attempts: joint acceptance accounts of group states (e.g. Gilbert (1989; 1990; 2014), Tuomela (1991; 1992)). Roughly, these accounts take these ascriptions of mental states at face value and accept that groups have mental states, such as beliefs and intentions. Moreover, joint acceptance accounts provide a particular story of how these group states come about in virtue of the group members jointly accepting such a group state. Being realist about these mental states and offering the joint acceptance story helps to explain not only our linguistic practice, but solves questions about a group's agency and the group's responsibility for action. It also provides an explanation for the commitments that group members take on board in virtue of jointly accepting something. However, joint acceptance accounts still do not cover all that ought to be said about the mental life of groups. One gap concerns the ability for groups to learn about their own mental states – group self-knowledge. The aim of this paper is to close this gap. I provide a list of desiderata that an account of group self-knowledge ought to explain and show how a theory of group self-knowledge for joint acceptance accounts can satisfy all these desiderata. As I will show in more detail, capturing group self-knowledge appropriately is important for any complete view of the mental life of groups, because it plays a role in the group's agency and moral responsibility for action. Therefore, it would be a virtue of joint acceptance accounts if they can be combined with a satisfying account of group self-knowledge. I aim to show that they can.

I start with a quick summary of joint acceptance accounts, before developing the desiderata for an account of group self-knowledge. I then propose my account and show how it satisfies all the desiderata.

2 Joint Acceptance Accounts

I begin with two assumptions. First, I am not providing a thorough discussion of what a group is. In our ordinary linguistic practice we refer to groups. Our governmental practices accept groups. Sociologists and mere common sense speak in favour of the existence of groups (Ritchie, 2015). In the following, I therefore presuppose that there are groups. Moreover, I presuppose that such groups come in different shapes and sizes. I take it that groups can be large, structured corporations, middle sized philosophy departments, small clubs and even just three people pushing a car together. All of these are groups in some sense. They are not just people arbitrarily taken together. There are different proposals on how groups differ from mere collections of individuals. The proposal most relevant for my purpose suggests that groups bind their members together in virtue of some form of shared or joint intention. In particular, joint acceptance accounts of group mental states commonly also take joint acceptance to make the difference between mere collection of individuals and proper groups. Groups are constituted by the members jointly accept something or jointly commit² to something (Tuomela, 2004; 2005; 2013 Gilbert, 1989; 1990; 2009; 2014).

Second, I am not going to defend the existence of mental states for groups. Joint acceptance accounts hold that there are such mental states and talk of them is not just metaphorical. Such group states are not uncontroversial (see for instance Quinton (1975) and Rupert (2005; 2019)), but it is not the aim of my paper to defend those states as such. Tollefsen (2002a) and Kallestrup (2020) point to evidence from ordinary linguistic and normative practice in which mental states of groups play a role. Tollefsen (2002a; 2002b) and Huebner (2014) argue for a group's mental states based on an intentional stance (cf. Dennett (1987)). Theiner et al. (2010)

² I will usually talk of 'joint acceptance' as the more general term including 'joint commitment'. There are small differences between those terms that will not be important for my purpose.

and Lyre (2018) point to socially extended knowledge as a path to mental states for groups. My entry into the debate is at a much later point: starting with joint acceptance accounts of group mental states, how should we understand group self-knowledge?

With these two assumptions in place, let me take a quick look at joint acceptance accounts. In its most general form these accounts accept some form of the following claim:

Joint Acceptance Account: A group G has attitude A iff the members of G jointly accept attitude A .

The details around this core can differ. For instance, the most prominent proposals by Gilbert (1989; 1990; 2014) and Tuomela (1991; 1992) also include a condition that it is common knowledge in G that the members have a willingness to jointly accept that p with other members of G .³ Following them, I also take such a common knowledge condition on board. Moreover, Tuomela makes the reasonable suggestion that not all group members matter, but only the subset that plays a role in a group's decision-making process. Call these *operative* members. This adjustment is required, because in many groups not every member has a say in a group's intentions or beliefs. For instance, the clerk at an Apple store will not influence what Apple intends to do as a group. A CEO on the other hand will. The notion of an operative member allows to capture this difference.

Some joint acceptance accounts (such as Gilbert's (1989)) also allow for non-basic cases of joint acceptance. These are cases in which the group members at one point jointly accept to grant individual members the power to form group states by themselves. For instance, the company members jointly accept that the finance officer is in charge of taxes and can form

³ Gilbert and Tuomela have slight differences in their details, but both have a condition that ensures that all relevant members are aware of the joint acceptance. In Tuomela this is not spelled out in terms of common knowledge, but rather as mutual knowledge about the joint acceptance.

group intentions related to taxes on their own. This can also be modelled in terms of operative members changing based on the subject matter of the group states in question.

Joint acceptance accounts are not the only option for theorizing about group mental states. Alternatives include interpretationist accounts by Tollefsen (2002a) or Huebner (2014), Searle's (1995; 2010) proposal of primitive we-attitudes located in individual agents, aggregation functions in List & Pettit (2011), and shared content views as proposed by Bratman (2014).⁴ While I will not provide a full defense of joint acceptance over these alternatives, I will mention two advantages of joint acceptance views discussed in the literature (e.g. Gilbert, 1989; 1990; 2009; 2014, Tuomela, 2004; 2005; 2013, Lackey 2020; 2021).

First, joint acceptance accounts are particularly well suited to explain the normative demands between group members. Gilbert (1990) explains this best with the example of walking together. Once two people agree to take a walk together, they are bound to certain commitments. Say Bob and Ross agreed to go for a walk. Bob starts walking way faster than Ross. It seems that Ross is entitled to call Bob out, telling him to walk slower. Gilbert explains this as a form of normativity established by the joint commitment (which is, roughly, Gilbert's version of joint acceptance). Many other accounts struggle with this sort of normativity. For instance, Bratman's (1993; 2014) shared content view with interlocking individual intentions cannot adequately explain the normativity involved in joint intention (Gilbert, 2009). Similarly, interpretationists, like Tollefsen (2002a), struggle with the normative demands on group members, because they tell us little about the role of group members in relation to one another and group mental states.

⁴ This list is not exhaustive.

Second, joint acceptance accounts are non-summative (or 'inflationist,' which I take to be synonymous) and enjoy some advantages because of that. They hold that mental states of groups cannot be understood merely in terms of all or most group members having a corresponding mental state. To illustrate, according to summative accounts a group believes that p iff all or most group members believe that p . Non-summatives, including joint acceptance accounts, reject that idea. They hold that groups can believe that p even if very few, perhaps even no group members at all believe that p and that sometimes a group might not believe p even though most members believe p . This has the upside of allowing non-summatives to explain paradigmatic divergence cases for group belief. Take for instance the following example:

Philosophy Department: The philosophy department at a leading university is deliberating about the final candidate to whom it will extend admission to its graduate programme. After hours of discussion, all of the members jointly agree that Jane Smith is the most qualified candidate from the pool of applicants. However, not a single member of the department actually believes this; instead, they all think that Jane Smith is the candidate who is most likely to be approved by the administration. (Lackey, 2020, p. 188)⁵

Summativists have a much more difficult time to explain how a group, such as the philosophy department, can form a belief in such a case. After all, the summative verdict would be that the department does not believe that Jane Smith is the most qualified candidate, given that no single member believes that. This gives the joint acceptance proponent a clear advantage.

⁵ Lackey herself is not a proponent of joint acceptance accounts, but the example is a good illustration of a divergence case. Similar cases have been proposed by Gilbert (1989), Schmitt (1994) and Tollefsen (2007).

The account can easily explain why the department believes Jane Smith is the most qualified: the group members jointly accept that she is.⁶

On the other hand, joint acceptance accounts also have their problems. I will just mention one significant worry. Jennifer Lackey (2020; 2021) argues that joint acceptance accounts give us the wrong verdict for cases of group lies. A tobacco company might be aware of scientific evidence linking smoking to cancer. However, the tobacco company's board members could nevertheless jointly accept that smoking is perfectly safe and healthy for humans. And they jointly accept this for financial reasons. When the company then claims in advertisement that smoking is perfectly healthy, it seems as though the group is telling a lie. However, the joint acceptance accounts would be committed to claiming that these advertisements do not include a lie at all. After all, the statement about smoking being healthy is what the group members jointly accept and therefore believe. Asserting a belief is not a lie, even if it might be a false assertion. This seems to force joint acceptance accounts into a wrong verdict of the tobacco case. Fortunately, the joint acceptance accounts have tools to explain what is going on in such a case. Bright (2020) and more recently Schwengerer (2022b) argue that joint acceptance accounts can explain why the tobacco company is blameworthy. They suggest that it is not a case of a group lie, but rather an epistemic mistake in the group's belief formation that we blame the group for. Hence, even with some difficulties the joint acceptance account of group states is a live option with overall explanatory benefits.

⁶ Joint acceptance accounts are not the only non-summativist option. For instance, some premise-based aggregation accounts (e.g. List & Pettit (2011)) can also be non-summativist. Premise based aggregation accounts are known to deal well with some divergence cases too, such as discursive dilemmas in which voting aggregation leads to group beliefs that come apart from the majority of group member beliefs. Joint acceptance theories can in principle also be combined with an aggregation of votes. This can be done either as a non-basic form of joint acceptance (e.g. when the group jointly accepted majority voting), or as a mechanism that motivates joint acceptance (e.g. when the group votes and the group members take the vote to be a reason in favor of jointly accepting the result of the vote).

Having given a short explanation of joint acceptance accounts I can now continue to the desiderata of group self-knowledge. In doing so I also show why group self-knowledge matters.

3 Desiderata for Group Self-Knowledge

To get one step closer to a full account of the mental life of groups based on a joint acceptance account I aim to provide a theory of self-knowledge within the joint acceptance framework. Such a theory has to capture the role that self-knowledge plays in our ordinary linguistic practice, a group's agency, and the responsibility a group has for its actions and mental states. And that role for group self-knowledge roughly matches the role that self-knowledge plays in individuals.

3.1 Our Linguistic Practice

The first desideratum is our ordinary linguistic practice involving self-ascriptions of mental states. Let me start by looking at individual self-knowledge and its relation to our linguistic practice. Generally, I am in an especially good position to tell anyone about my current mental states. If I say "I feel cold", then usually this will be accepted without doubt. Rarely anyone challenges such self-ascriptions, except perhaps by challenging my sincerity⁷. But if I am taken to be sincere, then it seems quite odd to challenge the truth of my claim that I feel cold. Individual self-knowledge seems to be *privileged*. Similarly, it seems rather inappropriate to ask for evidence for such a claim. "How do you know?" is usually not a felicitous response to my claim that I feel cold. I just know that I feel cold! Individual self-knowledge seems *peculiar* in its relation to evidence. How exactly these intuitions about privilege and peculiarity ought to be taken into consideration for an account of self-knowledge is controversial. Wright (1998;

⁷ Or perhaps by challenging my conceptual competences.

2015) and Finkelstein (2003) propose that what these intuitions show is primarily about our linguistic practice of avowing⁸. Many others take the intuitions to show something about the nature of second-order beliefs and their production (e.g. Armstrong (1968), Burge (1996), Moran (2001), Nichols & Stich (2003), Byrne (2005; 2018), Goldman (2006), Boyle (2009; 2011), Fernández (2013), Schwengerer (2019)).⁹ Regardless, any account of self-knowledge has to explain these features of our linguistic practice. This is a widely accepted desideratum for accounts of individual self-knowledge.

Given this desideratum for theories of individual self-knowledge one would expect a very similar desideratum for theories of group self-knowledge. Schwengerer (2022a) argues for such a desideratum at length. Just like individual avowals have a form of authority, so do avowals by groups. Only in rather unusual circumstances we would challenge group avowals, like mission statements, political manifestos, or press releases outside of their sincerity. Consider the manifesto case. Challenging political manifestos is commonplace, but the challenge is normally about the sincerity, or about how well the ideas in the manifesto will work to solve a particular problem. It would be rather odd to challenge a manifesto by claiming that, even though the party is sincere, what is written in it is not what the political party intends to do. This authority ascribed to group avowals is not to be read as a claim that it is *never* appropriate to challenge the avowal without questioning the sincerity. Rather, it should be read as a form of weak authority holding that such an appropriate challenge is *especially*

⁸ Usually understood as a speech act of self-ascribing a mental state. Explicit arguments against this focus on linguistic features have been provided by Borgoni (2019) and Schwengerer (2021).

⁹ However, Carruthers (2011) and Cassam (2014) suggest that this is merely an apparent privilege and peculiarity. According to them, self-knowledge is formed based on evidence and should be treated similar to other beliefs. Hence, they hold that there are more occasions in which avowals can and should be challenged appropriately than is often thought by other epistemologists. But they still agree that many everyday instances of me saying that I am in pain, or that I believe that it is raining seem to have an apparent authority that is reflected in our linguistic practice.

rare. This is also the case for the authority of individual's avowing their propositional attitudes (Wright, 1998, p. 17).

Schwengerer points to a second feature of group avowals that any account of group self-knowledge should explain: a weak form of *saliency*. Saliency is a term taken from Crispin Wright (2015) to capture the following observation about our linguistic practice:

If one is asked 'Do you believe that p?' it seems absurd to answer 'I don't know.' It is clearly not absurd to answer that you suspend judgment, but it would be absurd if you did not know whether you believe that p, not believe that p, or are withholding judgement. (Schwengerer, 2022a, p. 1156)

Saliency holds for avowals in individuals. For group avowals, saliency only holds with some qualifications. Call this *weak saliency*. Groups cannot always answer a question such as 'Do you believe that p?' immediately. In some cases, groups will not have considered whether p yet and take time to deliberate. In other cases, the group structure is such that the person (or persons) that can answer the question are not identical with those that are asked. Hence, producing the answer will take time until the internal group communication takes place. Moreover, sometimes a group can be asked 'Do you believe that p?' when it has not settled the question and is currently not even in a position to settle the question. For instance, some operative group members might be missing and therefore the group cannot deliberate whether p properly. To capture these qualifications weak saliency is understood as saliency only when the group properly received the question and the group is fully operational and given sufficient time. Importantly, that saliency here is qualified as 'weak' does not mean that the absurdity to answer 'I don't know' comes in any degree. Either it is absurd, or it is not. What the qualifier signifies is that there are more situations in which saying 'We do not know

whether we believe that p' is not absurd for groups, compared to individual agents. For individual agents, there are very few – if any – situations in which asserting ignorance of whether one believes p is appropriate. Wright (1998) formulates salience as though there are none, but even for individuals there might be special cases (e.g. heavy concussions) in which we would accept 'I don't know' as an answer to the question 'Do you believe that p ?'. Schwengerer (2022a) holds that for groups there are more situations in which 'I don't know' is acceptable and not obviously absurd, hence the claim of weak salience.

Following Schwengerer I take these two observations about the ordinary linguistic practice to be a desideratum of an account of group self-knowledge. Importantly, one feature of avowals in individuals is missing from group avowals. Asking for evidence is inappropriate in response to an individual's avowal. However, as Schwengerer (2022a) argues, this is not the case for group avowals. A group can appropriately be asked how the group knows that they believe p and the group can answer by, for instance, pointing to internal discussions, or the expertise of particular group members, or a vote that took place (Schwengerer, 2022a, p. 1162).

Based on this discussion I can formulate two desiderata for an account of group self-knowledge:

Desideratum 1: The account of group self-knowledge has to explain why group avowals are weakly authoritative.

Desideratum 2: The account of group self-knowledge has to explain why group avowals are weakly salient.

3.2 Practical and Moral Considerations

While linguistic features of avowals are a starting point for self-knowledge that is very widely accepted for theories of self-knowledge, there are some practical and moral considerations

that are not always at the center of attention. Nevertheless, these can be mobilized to show what function self-knowledge can have and why we should care about an account of group self-knowledge. This is well reflected in recent debates on the possibility of group self-knowledge. For instance, Jared Peterson states:

As noted above, inflationists hold that some groups act on the basis of their beliefs and desires. Such beliefs and desires, it is reasonable to believe, serve as an explanation for why groups do what they do. To return to an example referenced above, if the CIA launched its own social media accounts in part because it wanted to improve its favorability ratings with the public, and it believed it could do so by maintaining a social media presence, then in order to know why it was doing what it was doing, the CIA would, it seems, need to know that it had the relevant beliefs and desires. And if groups fail to know why they are behaving the way they are behaving, they would appear to be less morally responsible for their actions. (Peterson, 2022)

And similarly, Schwengerer (2022a) repeatedly points to practical and moral considerations to motivate why group self-knowledge matters. Take the following consideration about the role of self-knowledge in holding groups responsible:

For instance, take the court case against BP regarding the Deepwater Horizon oil spill. One of the crucial points of the case was the question whether BP was merely negligent, or whether it was a case of willful misconduct. That is, the question was precisely whether BP acted in a way that ignored relevant knowledge of the risks involved. One strategy to argue against the willful misconduct case would have been to claim ignorance of the risks. [...] BP might testify that they did not believe that there were any unusually high risks. How we should handle disagreement with BP's

testimony depends on whether we take BP to have any sort of privileged access to their intentional states. (Schwengerer, 2022a, p. 1158)

These discussions present a clear connection between self-knowledge, agency, and responsibility for action. Following a plausible picture of action, groups act because of their beliefs and desires. And access to their own beliefs and desires allows for a form of control and guidance in their action. If they do not know what they want to achieve, nor what the group believes, then it seems difficult for the group to coordinate the actions of group members such that they can achieve the group's goal. Moreover, as Peterson (2022) rightly mentions, in order to answer questions about the reasons for their actions they need self-knowledge about their beliefs and desires. And answering the question why the group wants to do something is important for group deliberations about the appropriate means. For instance, a group might desire more funds and deliberate on how to acquire such funds, given their beliefs about their current situation. What kind of action the group considers as an available option will depend on the group's attitude regarding their own desire. If they take more funds to be a question of life and death for the group they might consider rather extreme means – perhaps even illegal means. On the other hand, if they don't take the desire to be so important – perhaps because they believe that they can also function without these funds – then the potential means will be different. The group deliberation will likely not consider any illegal ways of acquiring funds in this case. But this difference in their deliberation can only be explained if the group knows their desire and beliefs. If the group does not know that they desire more funds it seems impossible for the group to start deliberating about the appropriate means to acquire those funds. It would make no sense to start such a deliberation. Deliberating on group actions seems only possible if the group can be aware of the group's

desires and beliefs. A group's critical evaluation of their own potential actions requires self-knowledge.

Closely connected to the evaluation of potential action is the group's critical evaluation of their own mental states. And that seems to relate to self-knowledge as well. This is a common idea from discussions of individual self-knowledge. The most prominent proponent of self-knowledge as a necessary condition for critical assessment of one's own mental states is Tyler Burge (1996). He takes a very demanding, rationalist conception of self-knowledge to be necessary to explain critical reasoning about mental states.¹⁰ I do not want to commit to any such conception. Instead I follow Peterson (2022) for a much weaker claim: having self-knowledge enables one to critically reason about one's mental states. There might be other ways to critically assess and reason about one's mental states. And these other ways might also be able to change one's mental states. But in any case, self-knowledge can bring a person into a position in which they can critically evaluate and reason about their mental state. And the same goes for groups. A group that is aware of its beliefs can assess the beliefs based on the available evidence and revise them if necessary.

As seen self-knowledge plays a role in critically evaluating both group actions and group mental states. This is something that should also be captured in an account of group self-knowledge. Moreover, both are closely connected with practical and moral considerations. Take Schwengerer's court scenario about BP. In such a case the court wants to find out whether BP knew about the risks involved in drilling for oil. The court wants to know whether BP had a particular belief, or a particular piece of knowledge. Whether BP had these mental states matters insofar as we hold BP responsible in different ways depending on whether BP

¹⁰ Which has been challenged by Gertler (2018).

knew, or whether BP did not know. But the responsibility question does not end here. Judgments about the responsibility of BP also look back at whether BP engaged in critical evaluation of their own states. It makes a difference whether BP carefully and critically considered all their knowledge and evidence about the risks, or whether BP did not do that at all. And an account of self-knowledge should tell us how easy it would have been for BP to critically reflect on the knowledge BP actually had. If it turns out that BP knew, but it was very difficult for BP to recognize that they knew and to adjust their actions based on that knowledge, then BP might be less responsible for the oil spill. On the other hand, if it turns out that BP knew and that it was very easy for BP to recognize that they knew, but they still did not bother to deliberate whether they know, nor to deliberate about potential consequences of such knowledge, then BP seems much more responsible for the oil spill. The details of the interactions between first-order knowledge, self-knowledge of that first-order knowledge, and agency matter. An account of group self-knowledge should tell us how easy or difficult it is for the group to recognize their mental states, to critically evaluate them, and to adjust their actions properly to those states. This then gives us also an idea of how much responsibility we should ascribe to groups for their actions and their mental states. Of course, self-knowledge is not the only thing that matters for responsibility. My aim is much more modest. Self-knowledge seems to be a contributing factor for the responsibility that groups have for their mental states and actions and this contribution of self-knowledge should be captured by my theory of group self-knowledge. An account of group self-knowledge ought to explain how self-knowledge influences this responsibility.

With this discussion in mind I end up with three more desiderata:

Desideratum 3: The account of group self-knowledge has to explain how groups can critically evaluate and reason about their potential actions.

Desideratum 4: The account of group self-knowledge has to explain how groups can critically evaluate and reason about their mental states.

Desideratum 5: The account of group self-knowledge has to contribute some part to an explanation of the extent to which groups are responsible for their actions and mental states.

I am sure that these desiderata are not exhaustive, but nevertheless they seem to be the most pressing puzzles an account of group self-knowledge should solve. And this is what I am going to do in the next section.

4 Self-Knowledge in a Joint Acceptance Account

Joint acceptance accounts have a clear story how first-order mental states are formed by joint acceptance of operative group members. To stay within this framework self-knowledge also has to be formed by joint acceptance. For instance, the operative group members have to jointly accept that the group believes that p .

This unique starting point for group self-knowledge in the joint acceptance framework rules out a straightforward application of self-knowledge theories for individuals to the group case. Joint acceptance usually happens under the control of the group members and is based on grounds available to those members. Many common theories of self-knowledge, however, take self-knowledge to be rather automatic, sometimes rejecting citable grounds for self-ascribing a mental state. Some propose a perception like inner-sense (e.g. Nichols & Stich (2003); Goldman (2006)), others explain self-knowledge as inferences from sensory data (Carruthers, 2011), automatic detection of states in conscious awareness (Macdonald, 2014),

or as a constitutive part of first-order belief (Shoemaker, 1996).¹¹ None of them seem a good fit for the joint acceptance framework. The closest analogue will be a form of maker's knowledge or an agentalist account based on making up one's mind (Moran, 2001). It is no perfect analogue, but the same procedures that make up a group's mind will also play a role in gaining group self-knowledge. However, given the joint acceptance framework, these procedures can only play a role as feeding into a joint acceptance. The joint acceptance proponent holds that the second-order belief – like any other group belief – is formed by a joint acceptance of the group members. I aim to find a way to establish which joint acceptances on what grounds can provide group self-knowledge. My answer is going to be that only grounds that are adequately connected to the formation of first-order group mental states will do. In this sense, it is a form of maker's knowledge.

My task now is the search for a good basis for a joint acceptance that amounts to group self-knowledge. If group members randomly decide to jointly accept that the group believes that *p*, then that does not seem good enough for self-knowledge. One could not explain any sort of authority of avowals, nor any connection to critically evaluating mental states – after all they might not even have the corresponding first-order state. What I need is some connection between the joint acceptance that the group believes that *p*, and the group's belief that *p*. Jointly accepting that the group believes that *p* ought not to be disconnected from the first-order belief the group has. Otherwise group self-knowledge would turn out to be a matter of luck. So, what I am looking for is a connection between joint acceptance that *p* and joint acceptance that the group believes *p*. This connection then allows me to explain the desiderata.

¹¹ The list is not exhaustive.

The first step in my search for such a connection is to look at a basing relation that works for joint acceptance accounts. A suitable version has recently been developed by Jessica Brown (2022) who provides a way to model how groups can believe for a reason that is compatible with joint acceptance accounts. This gives me the desired connection by explaining how joint acceptance that the group believes that *p* can be based on a reason that is sufficiently related to the group's first-order belief. This is exactly the connection between first-order belief and second-order belief that I am looking for. Brown proposes a causal account of the basing relation for groups:

Basing: A group believes that *p* on the basis of reason *r* if and only if the group's belief that *p* is non-deviantly caused by *r*.

The details of how to spell out 'non-deviantly' are tricky and not important for my purpose.¹² What is important is the idea that Brown's account is neutral about the nature of *r* and how *r* exactly is established. The general formulation is compatible with *r* being generated by group deliberation ending up in joint acceptance, by some form of voting, or by division of labor with only a subset of the group establishing *r*. All that matters is that there is a causal link between *r* and the group belief – whatever *r* is and however it came to be. The account also allows for divergence cases similar to the one discussed earlier. Group members might believe that *p* because of *r*₁, but the group belief itself could still be caused by a different reason *r*₂ (Brown, 2022). It is helpful to fully translate **Basing** into joint acceptance terms:

¹² Brown is sympathetic to understanding non-deviantly caused by an appeal to dispositions. She points to Turri (2011, p. 393) who suggests that 'R non-deviantly causes B if and only if R's causing B manifests (at least some of) your cognitive traits'. I share Browns' sympathy for this understanding of non-deviant causes as it successfully rules out causality routes that are not sufficiently related to one's cognition.

Basing-JAA: A group believes that *p* on the basis of reason *r* if and only if the operative member's joint acceptance that *p* is non-deviantly caused by *r*.

It seems clear in this formulation that the causal story has to go through the operative members. Hence, some *r* has to cause the group members to jointly accept *p*.

To use this basing relation in my account of group self-knowledge I need to establish how the first-order group belief can be either itself the reason for the group self-knowledge, or can at least be linked to a reason for the group to believe that the group believes *p*. The model has to hold that a group believes that the group believes *p* on the basis of reason *r* and that reason *r* is at least closely linked to the group's belief that *p*. This link has to be such that the group's second-order belief based on *r* is often enough true, such that the process of generating group self-knowledge is reliable. Importantly, *r* does not need to be the first-order group belief itself. It merely needs to be sufficiently linked to that belief.

There are at least three options that might be good candidates for a reason *r* that can be the basis of group self-knowledge of the group's belief that *p*: the group's belief that *p* itself; or the individual operative group members' beliefs that the group believes *p*; or the individual operative group members' beliefs that they jointly accepted *p*.

It turns out that the best option is the group members' belief that they jointly accepted *p*. The argument here is twofold. First, the group belief that *p* seems to entail that the group members believe that they jointly accepted *p*. The argument for this has been provided by Peterson (2022) in a different context. Peterson argues that in general there cannot be any cases in which a group has a belief that *p*, without any group members knowing that they have jointly accepted *p*. In order for a group to form the belief that *p*, the group members need to jointly accept *p* under conditions of common knowledge. The condition of common

knowledge then entails that every group member that is part of this joint acceptance knows that they are jointly accepting p. In Peterson's words:

Consider: if a group comes to jointly accept that P, and no members of the group believe that P, then it is at least somewhat plausible to think that the group has a belief that the members don't. However, it is implausible to think that the members of such a group would fail to have the belief that they just jointly accepted the proposition in question. After all, they just went through the conscious and deliberate process of accepting that very proposition. Such members are, therefore, not only going to be aware of having gone through such a process, but they are also going to believe that the group has just jointly accepted the proposition in question. (Peterson, 2022, p. 16)

There does not seem to be any room for a situation in which the group believes p without the group members believing that they jointly accepted p leading to the group's belief that p. Such a case seems impossible in a joint acceptance framework. This gives us the first part of the argument for the members' beliefs that they jointly accepted p as the best basis for group self-knowledge of the group belief that p. If this is right, then there is no reason to take the group belief that p to be the basis for group self-knowledge, instead of the group members' beliefs that they jointly accepted p. Both coincide, but it is much easier to make sense of the causal efficacy of the beliefs located in the group members.

Second, the members' belief that they jointly accepted p is also a better choice than the members' belief that the group believes p. In many cases group members will have both. However, it seems possible for group members to have the former without the latter. A member might be part of a joint acceptance without thinking about any consequences of that

joint acceptance for the group belief. Hence, that member lacks a belief about the group belief. They probably could easily form such a belief, but they just have not yet done so.

I end up with the group members' beliefs that they jointly accepted *p* as the best choice for a basis of group self-knowledge. With that in mind group self-knowledge of belief can be accounted for as follows:

Group Self-Knowledge for Belief

A group has self-knowledge of its own belief that *p* if and only if

- (i) the operative members jointly accept that the group believes *p*; and
- (ii) that joint acceptance is non-deviantly caused by *the operative group members' beliefs that the members jointly accepted p* (the reason *r*); and
- (iii) the group believes *p*.

Condition (i) captures that the group forms the relevant second-order belief. Condition (ii) captures the proper basing relation that connects the second-order belief to the first-order belief. Finally, condition (iii) is a truth condition to transform the second-order belief into knowledge.

I can also formulate a more general version:

Group Self-Knowledge for Attitudes

A group has self-knowledge of its own attitude *A* if and only if

- (i) the operative members jointly accept that the group has attitude *A*; and
- (ii) that joint acceptance is non-deviantly caused by *the operative group members' beliefs that the members jointly accepted attitude A* (the reason *r*); and
- (iii) the group has attitude *A*.

The connection between first-order attitude and second-order belief in (ii) provides justification for the second-order belief. The second-order belief has a proper basis that is a reliable indicator for the first-order attitude. That is, usually when the operative group members' beliefs that the members jointly accepted A causes the group members to jointly accept that the group has attitude A, it will be the case that the group has attitude A.

Let me illustrate this with an example. Suppose the board of directors of Apple come together and jointly accept that they want to release a new iPhone. Each member expresses openly a willingness to jointly accept as such. This establishes a joint intention and hence Apples' intention. If the group now wants to find out what they intend to do, the group members together can make that happen by considering their joint acceptance. Every group member knows that they have jointly accepted to release a new iPhone. Because they all have this knowledge, they are now willing to take on another joint acceptance: the joint acceptance that the group intends to release a new iPhone.

In practice the group members might talk and discuss about what they already jointly accepted. In this discussion, the operative group members will quickly agree that they jointly accepted to release a new iPhone. After all, this is something that every single operative member already believes. Given that, the operative group members can (and likely will) express openly that they are willing to jointly accept that the group intends to release a new iPhone. When all operative group members have expressed that willingness to jointly accept that the group has that intention, they thereby form a second-order group belief about their intention. Their second-order belief will be true, because the group has the relevant first-order intention. And the second-order belief will be justified, because it is based on the group members' beliefs that they jointly accepted to release a new iPhone which reliably indicates the group intention. Hence, the group knows its own intention.

This reliable indication is important and worth looking into more closely. Peterson's argument shows that in any case in which the group believes p by joint acceptance, the operative group members will be aware of that joint acceptance. This establishes that whenever the group has a belief, the group members also have a reason r that can be the basis for a new joint acceptance that the group believes p . This ensures reliability in one direction: if the group believes p , then the members have a reason r . To fully establish that r is a reliable indicator for the group belief that p , one also has to show that r is unlikely to show up without the group having a belief that p . Can the reason r in (ii) be present without the group believing p ? Can all operative group members believe that they jointly accepted p , even when they did not?

This does not seem to be easy. I can imagine without a problem that some individual group members mistakenly believe that all group members jointly accepted p . But that will not be enough to be a worry for group self-belief, because in such a case the group members will not jointly accept that the group believes p .¹³ Some other group members will refuse to form a joint acceptance that the group believes p , because they do not believe that they jointly accepted p . Hence, (i) would not be satisfied.

But what if all members mistakenly believe that they jointly accepted p ? It is difficult to imagine such a scenario. After all, if they all believe they jointly accepted p , the group will act as if p . That is, the group will believe p . So, the group did jointly accept that p . No mistakes to be found. There does not seem to be room for all operative members to believe that they jointly accept p , without the group actually having the belief that p . Hence, whenever the operative group members have the reason r that can be the basis for a new joint acceptance that the group believes p , the group also believes p . This ensures reliability in the other

¹³ Nor will common knowledge conditions be satisfied.

direction: if the operative members have reason r , then the group believes p . Both directions combined give us a good reason to believe that r is a reliable indicator for the group believing p . It is such a reliable indicator that it ensures both justification and an anti-luck condition. In almost all (or all) nearby close possible worlds in which the second-order belief is formed based on r , the second-order belief will be true.

This account of group self-knowledge gives me everything I need to capture successful instances of group self-knowledge and to explain failures of group self-knowledge. In successful cases the operative members' joint acceptance is based on an appropriate reason r that is a reliable indicator for the first-order group mental state. In a failing case the group members jointly accept that the group has, for instance, a belief, but do not base that joint acceptance on an appropriate reason. Their joint acceptance is unrelated to the first-order group belief, and hence likely to be false, or merely true by luck.

My account is non-summative in the following sense: it denies that member belief that the group has attitude A is sufficient for group self-belief (group self-knowledge in the good case). It nevertheless is compatible with the idea that member belief is necessary for group self-belief (group self-knowledge). This is important, because it helps me to avoid Peterson's (2022) arguments against non-summative, inflationary group self-knowledge. Peterson argues that a divergence of group self-knowledge from member knowledge about the group mental states is impossible. He starts in the passage I cited above, in which he shows that whenever a group belief is established by joint acceptance the group members have beliefs about that joint acceptance. This part I am fully on board with.

The passage from above then continues as follows: "Ergo, there is not going to be a divergence between what the group believes about the target proposition concerning the group mind

and what the individuals believe” (Peterson, 2022, p. 16). Here two things are taking place. First, Peterson transitions from group member belief about joint acceptance to group member belief about “the target proposition concerning the group mind”. These are importantly different and I take this to mean that he is now looking at the group member belief about the group belief. I am not fully on board with this transition. A group member might have the belief about joint acceptance without having the belief about the group belief. I might be aware of the joint acceptance, but in that moment just not think about what that means for group beliefs. Nevertheless, I can grant Peterson that it would not take much for a group member who believes that they jointly accepted that *p*, to also believe that the group believes *p*. So, let me assume for the argument that Peterson indeed has shown that group belief always comes with group member beliefs about the group belief.

Second, Peterson now wants to show that divergence of group self-knowledge and member belief about the group mental state is impossible. I have granted that whenever a group belief is established by joint acceptance, the group members know of that group belief. But that does not show us that “there is not going to be a divergence between what the group believes about the target proposition concerning the group mind and what the individuals believe” (Peterson, 2022, p. 16). At most this establishes that whenever there is group self-knowledge, group members also have the corresponding belief about the group mental state. Group member beliefs about the group mental state turn out to be necessary for group self-knowledge. But that does not rule out all forms of divergence, because, as my account suggests, group member beliefs about group mental states are not sufficient. Group self-knowledge needs a distinct joint acceptance that the group has a particular mental state, as captured in condition (i).

In fairness to Peterson, his main target are divergence cases in which group self-knowledge happens without any group members knowing about the group's mental states. But he overshoots the target with the claim that there are not going to be any divergences between group self-knowledge and member knowledge about the group mental states. My proposal predicts – rightfully I think – some possible divergence cases in which group members believe that the group believes p , but the group lacks self-knowledge of the belief that p .

My resulting proposal is unique to the joint acceptance framework. Given that the basis for group self-knowledge is the group member's beliefs about joint acceptance, no other view of group mental states can make use of my proposal. This might turn out to be an advantage for joint acceptance accounts in general. If my proposal is right, then joint acceptance accounts have access to a good explanation of the features that group self-knowledge and its linguistic expressions have. Other accounts of group mental states have to bring something similar to the table. Currently, for other views of group mental states there is nothing on offer other than a hint at a transparency approach in Schwengerer (2022a) and an inferentialist account (Schwengerer, 2023a). This inferentialist account suggests that groups can observe the group's behavior and deliberately generate evidence that they can use to infer their group mental states by prompting discussions within the group about hypothetical scenarios. For instance, to find out whether Microsoft desires to expand, they can consider and discuss what they would do if competitors could be potentially bought and taken over. If in many hypothetical scenarios they conclude that they would buy the competitor, that would be considered sufficient evidence for the group having a desire to expand. The idea is that this behavior in hypothetical scenarios would be best explained as being caused by that desire to

expand.¹⁴ Such an inferentialist approach might work at times – in particular for inferring a group’s vices, as is the aim in Schwengerer (2023a). However, it struggles to give a plausible account that satisfies the desiderata in ordinary cases of knowing group attitudes. It requires a very long-winded process of internal group discussions that take time before the evidence generated can be used for a sufficiently reliable inference. In order to explain any authority of group self-knowledge, the inferential account needs such a reliable inference. Hence, this inferentialist account has problems explaining how most of the group’s avowals are weakly authoritative, given that most of them are not based on extensive internal discussions that could generate sufficient evidence in the inferentialist picture. This is not a problem for the aims of Schwengerer (2023a), who wants to explain self-knowledge of group vices, and self-knowledge of vices is usually taken to be far less – if at all – privileged compared to self-knowledge of mental states. But it shows that the inferentialist account is not well-suited to explain general group self-knowledge.

Let me now show how my current proposal can do better and satisfies the identified desiderata.

4.1 Satisfying the Desiderata

The account should satisfy the following five desiderata:

Desideratum 1: The account of group self-knowledge has to explain why group avowals are weakly authoritative.

Desideratum 2: The account of group self-knowledge has to explain why group avowals are weakly salient.

¹⁴ This is similar to the inferentialist accounts of Lawlor (2009) and Cassam (2014) in individuals, that explain the acquisition of self-knowledge as an inference to the likely cause of behavior, thoughts, feelings, etc.

Desideratum 3: The account of group self-knowledge has to explain how groups can critically evaluate and reason about their potential actions.

Desideratum 4: The account of group self-knowledge has to explain how groups can critically evaluate and reason about their mental states.

Desideratum 5: The account of group self-knowledge has to contribute some part to an explanation of the extent to which groups are responsible for their actions and mental states.

The first two desiderata concern our ordinary linguistic practice. Groups can avow their mental states and these avowals are weakly authoritative and weakly salient. In discussions of individual self-knowledge authority and salience are often explained in virtue of an especially easy and reliable process of acquiring self-knowledge. Suppose it is very easy to know one's own mental state, and one only very rarely is mistaken about one's own mental state. If this was the case, then one has a natural explanation for both authority and salience. Avowals appear to be authoritative because self-beliefs are usually true. Challenging my avowal is usually inappropriate because it is so unlikely that I am wrong about myself.¹⁵ Moreover, avowals appear salient, because self-beliefs are very easy to acquire and not having them is unusual. If I am asked 'Are you in pain?' it would seem odd to say 'I don't know' because knowing whether I am in pain is so easy. Something very unusual has to be going on in order for me not to be able to say whether I am in pain. I can now provide a very similar story for group self-knowledge.

¹⁵ Of course, such an explanation is also contested for individual self-knowledge. See for instance Schwitzgebel (2008), Carruthers (2011), or Cassam (2014).

Group self-knowledge is very easy to come by in my proposed model. Usually when a group believes p , the group members are aware of them having jointly accepted p . And they can use that awareness as a reason to jointly accept that the group believes p . As Peterson shows, it seems impossible that a group believes p in virtue of joint acceptance, without the group members being aware of that joint acceptance. So, all that is required for the group to acquire self-knowledge is that the group members act on their available reason to jointly accept that the group believes p . Of course, they might not form this further joint acceptance. The group might stay without self-knowledge, if they do not see any reason to deliberate on their own belief. But if the group is interested in their own belief, the group has an easy way to acquire self-knowledge. And that ease is enough to explain the linguistic features of group avowals. When a group sincerely avows that they believe p , they do so because the group believes that they believe p . And given that the group is in an especially good position to form that second-order belief, a challenge seems usually misplaced. This is exactly the weak authority I want to capture.

And for very similar reason weak salience is explained by my account. It seems odd for a well-functioning group to answer the question 'Do you believe p ?' with 'We don't know'. It seems odd and inappropriate because it is very easy for a well-functioning group to answer that question. If the group believes p , and the group is currently functioning properly (e.g. all operative members are present), then they are in a position to know whether they believe that p . They have a reason available that allows them to jointly accept that the group believes p . A similar story can be given if the group believes non- p , or if the group believes neither. The group will be in a position to tell, because the group members are aware of the joint acceptance that grounds their first-order belief.

The desiderata 3 and 4 relate to the critical assessment and reasoning about actions and mental states. Self-knowledge enables one to critically evaluate and potentially change one's actions and mental states. Group self-knowledge should be able to play the same role for groups. And my proposed account allows this to a high degree. Given that any time a group forms a belief that p the group members are aware of the joint acceptance that forms the belief, the group members always have a reason r that can be used in joint acceptance to form group self-knowledge. Hence, a group believing p is always in a position to know of this belief. This allows for the group to access their own beliefs rather easily and opens those beliefs up for critical evaluation within the group. Groups can easily re-evaluate the beliefs previously formed by joint acceptance. And the operative group-members can change the beliefs based on that re-evaluation by performing another act of joint acceptance.

Easy self-knowledge of beliefs and desires also allows the group to critically evaluate its actions based on those beliefs and desires. According to my account, a group can very easily know what the group desires. This allows the group to reflect on that desire, and to pick means to fulfil the desires that are adequate for the situation.

For both critical reflection of mental states and actions it is important to emphasize that the ease of self-knowledge in my model does not provide any sort of guarantee. They group can easily know its beliefs and desires. But it is still a matter of whether the group wants to know them. Group self-knowledge requires that the operative group members form a joint acceptance that the group (for instance) believes, or desires something. But if the group members do not form such a joint acceptance, then the group will not have self-knowledge. There is no form of self-intimation for group mental states that grants group self-knowledge for free. Group self-knowledge will always require the group members to work at least a little

bit, and there will always be the danger that the group members refuse to do that.¹⁶ Moreover, there is also room in the model for groups to form beliefs about their own mental states for the wrong reason. Self-knowledge is only acquired if the operative members jointly accept that the group has a mental state for a reason that is linked to that state. If they jointly accept for any other reason, they can easily be wrong about themselves. This leaves room open for groups to deceive themselves by turning a blind eye to their first-order joint acceptances.

The final desideratum concerns the role of self-knowledge in the responsibility that groups have for their actions and mental states. Based on what I have already said, it will not be surprising that the easy access to group self-knowledge leads to ascribe responsibility to the groups. Self-knowledge allows groups to know why they act the way they do. And such knowledge adds to the group's responsibility for that action (Peterson, 2022). It allows the group to rethink its actions and change, adjust or simply stop acting if they take the reasons for why they act to be insufficient to justify their action. If the group nevertheless continues to act for their reason, they thereby seem to endorse the action and the reason for their action. They seem to endorse action and reason precisely because self-knowledge comes easy to the group. The group has no excuse for not reflecting on the reasons for their actions. This explains why we hold groups responsible for their action in part because of the group's self-knowledge. Given that we want to hold groups responsible for their actions, and that we already do so (e.g. in court cases), this is a welcome result.

The account can provide a further avenue for holding groups responsible beyond the knowledge of the group's reasons for actions. Goldberg (2016; 2018) has proposed that

¹⁶ Whether a particular group tends to be more or less reflective about their own states and actions could be analyzed in terms of epistemic virtues and vices of a group. I will bracket this issue here.

people have epistemic expectations towards other agents. For instance, I expect my doctor to be aware of side effects of medication and that they are reasonably informed about the latest studies on the medication they prescribe. I suggest that there are also epistemic expectations in regard to agents critically evaluating their actions and mental states. People tend to blame others for not thinking their actions through, or for blindly believing something they read. This is a form of blame given in response to an agent that has not critically evaluated their actions and mental states. And the same sort of blame can be attributed to groups. It seems appropriate to blame a group for not having properly deliberated on the consequences of their actions, and whether those consequences were justified with respect to the group's goals. The group did not think things through. And this can be captured in terms of the group failing to satisfy an epistemic expectation. My account of self-knowledge provides an idea of how these epistemic expectations towards a group can arise. These expectations are a result of the easily available path to self-knowledge that groups have. We expect the group to critically evaluate its actions and mental states at least in part because it is so easy to do for the group.¹⁷ It is not difficult for the group to know what it wants to do, and why it wants to do it. If the group nevertheless does not critically reflect on its actions and mental states, then the group is acting intellectually careless. This is what we blame them for. And this explanation fits well within my account of group self-knowledge. Overall it seems that my account can explain how group self-knowledge contributes to the responsibility that groups have for their actions and their mental states. The account satisfies all desiderata and is therefore a good candidate for a theory of group self-knowledge.

¹⁷ For more discussion of epistemic expectations and easy availability see Schwengerer (2023b).

5 Conclusion

I have provided an account of group self-knowledge in the joint acceptance framework. In general, a group has self-knowledge of its own attitude A if and only if (i) the operative members jointly accept that the group has attitude A; and (ii) that joint acceptance is non-deviantly caused by the operative group members' belief that the members jointly accepted attitude A (reason r); and (iii) the group has attitude A. Together (i) to (iii) give us conditions for a group having a justified, non-lucky, true belief about its own attitude.

The proposed account satisfies all desiderata I have identified. It can explain why group avowals are weakly authoritative and weakly salient. It can also explain how group self-knowledge enables the group to critically evaluate their actions and mental states. And finally, it can explain how group self-knowledge contributes to the group's responsibility for its actions and mental states. I am certain that some of these aspects can be looked at in even more detail. However, the current proposal is a good step in the right direction.

6 Bibliography

Armstrong, D. M. (1968). *A Materialist Theory of Mind*. London: Routledge & Kegan Paul.

Borgoni, C. (2019). Authority and Attribution: the Case of Epistemic Injustice in Self-Knowledge. *Philosophia*, 47(2), 293-301.

Boyle, M. (2009). Two Kinds of Self-Knowledge. *Philosophy and Phenomenological Research*, LXXVIII(1), pp. 133-164.

Boyle, M. (2011). Transparent Self-Knowledge. *Proceedings of the Aristotelian Society Supplementary*, LXXXV, pp. 223-241.

Bratman, M. E. (1993). Shared Intention. *Ethics*, 104(1), 97-113.

- Bratman, M. E. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Bright, L. K. (2020). Group Lies and Reflections on the Purpose of Social Epistemology. *Aristotelian Society Supplementary Volume*, 94(1), pp. 209-224.
- Brown, J. (2022). Group Belief for a Reason. *Aristotelian Society Supplementary Volume*, XCVI.
- Burge, T. (1996). Our Entitlement to Self-Knowledge. *Proceedings of the Aristotelian Society, New Series*, 96, pp. 91-116.
- Byrne, A. (2005). Introspection. *Philosophical Topics*, 33(1), pp. 79-104.
- Byrne, A. (2018). *Transparency and Self-Knowledge*. Oxford: Oxford University Press.
- Carruthers, P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. Oxford: Oxford University Press.
- Cassam, Q. (2014). *Self-Knowledge for Humans*. Oxford: Oxford University Press.
- Dennett, D. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.
- Fernández, J. (2013). *Transparent Minds: A Study of Self-Knowledge*. Oxford: Oxford University Press.
- Finkelstein, D. (2003). *Expression and the Inner*. Cambridge, MA: Harvard University Press.
- Gertler, B. (2018). Self-Knowledge and Rational Agency: A Defense of Empiricism. 96(1), pp. 91-109.
- 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.

- Gilbert, M. (2009). Shared Intention and Personal Intention. *Philosophical Studies*, 144(1), 167-187.
- Gilbert, M. (2014). *Joint Commitment*. Oxford: Oxford University Press.
- 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.
- Goldman, A. (2006). *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford: Oxford University Press.
- Huebner, B. (2014). *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford: Oxford University Press.
- Kallestrup, J. (2020). Group virtue epistemology. *Synthese*, 197, 5233–5251.
- Lackey, J. (2020). Group Belief: Lessons from Lies and Bullshit. *Aristotelian Society Supplementary Volume*, 94(1), pp. 185-208.
- Lackey, J. (2021). *The Epistemology of Groups*. Oxford: Oxford University Press.
- List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Lyre, H. (2018). Socially Extended Cognition and Shared Intentionality. *Frontiers in Psychology*, 9. doi:10.3389/fpsyg.2018.00831
- Macdonald, C. (2014). In my 'Mind's Eye': introspectionism, detectivism, and the basis of authoritative self-knowledge. *Synthese*, 191(15), pp. 3685-3710.

- Moran, R. (2001). *Authority and estrangement*. Princeton: Princeton University Press.
- Nichols, S., & Stich, S. (2003). *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding Other Minds*. Oxford: Oxford University Press.
- Peterson, J. (2022). Groups that fly blind. *Synthese*, 200(6), pp. 1-24.
- Quinton, A. (1975). Social Objects. *Proceedings of the Aristotelian Society*, 75, 67-87.
- Ritchie, K. (2015). The Metaphysics of Social Groups. *Philosophy Compass*, 10(5), pp. 310-321.
- Rupert, R. (2005). Minding One's Cognitive Systems: When Does a Group of Minds constitute a Single Cognitive Unit? *Episteme*, 1(3), pp. 177-188.
- Rupert, R. (2019). Group Minds and Natural Kinds. *Avant: Trends in Interdisciplinary Studies*, 10(3). doi:10.26913/avant.2019.03.08
- 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. (2019). Self-Knowledge in a Predictive Processing Framework. *Review of Philosophy and Psychology*, 10(3), 563-585.
- Schwengerer, L. (2021). Beliefs Over Avowals: Setting Up The Discourse On Self-Knowledge. *Episteme*, 18(1), pp. 66-81. doi:10.1017/epi.2018.56
- Schwengerer, L. (2022a). Towards Collective Self-Knowledge. *Erkenntnis*, 87(3), pp. 1153-1173.
- Schwengerer, L. (2022b). Defending Joint Acceptance Accounts of Belief Against the Challenge from Group Lies. *Logos and Episteme*, XIII(4), pp. 421-428.

- Schwengerer, L. (2023a). Collective vice and collective self-knowledge. *Synthese*, 201(19).
- Schwengerer, L. (2023b). An Anchored Joint Acceptance Account of Group Justification. *Theoria*, 89(4), pp. 432-450.
- Schwitzgebel, E. (2008). The Unreliability of Naive Introspection. *Philosophical Review*, 117(2), pp. 245-273.
- Searle, J. (1995). *The Construction of Social Reality*. New York: The Free Press.
- Searle, J. (2010). *Making the Social World: The Structure of Human Civilization*. New York: Oxford University Press.
- Shoemaker, S. (1996). *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing Group Cognition. *Cognitive Systems Research*, 11(4), pp. 378-395.
- Tollefsen, D. (2002a). Organizations as True Believers. *Journal of Social Philosophy*, 33(3), pp. 395-410.
- Tollefsen, D. (2002b). Collective Intentionality and the Social Sciences. *Philosophy of the Social Sciences*, 32(1), pp. 25-50.
- Tollefsen, D. (2007). Group Testimony. *Social Epistemology*, 21(3), 299-311.
- Tuomela, R. (1991). We Will Do It: An Analysis of Group-Intentions. *Philosophy and Phenomenological Research*, 51(2), pp. 249-277.
- Tuomela, R. (1992). Group Belief. *Synthese*, 91(3), pp. 285–318.
- Tuomela, R. (2004). Group Knowledge Analyzed. *Episteme*, 1(2), 109-127.

Tuomela, R. (2005). We-intentions Revisited. *Philosophical Studies*, 125(3), 327–369.

Turri, J. (2011). Believing for a Reason. *Erkenntnis*, 74(3), pp. 383-397.

Wright, C. (1998). Self-Knowledge: The Wittgensteinian Legacy. In C. Wright, B. Smith, & C. Macdonald, *Knowing our Minds* (pp. 13-45). Oxford: Oxford University Press.

Wright, C. (2015). Self-knowledge: the reality of privileged access. In S. C. Goldberg (Ed.), *Externalism, Self-Knowledge, and Skepticism* (pp. 49-74). Cambridge: Cambridge University Press.