



Group belief reconceived

Jeroen de Ridder¹

Received: 8 February 2021 / Accepted: 21 February 2022 / Published online: 15 April 2022
© The Author(s) 2022

Abstract

An influential account of group belief analyzes it as a form of *joint commitment* by group members. In spite of its popularity, the account faces daunting objections. I consider and reply to two of them. The first, due to Jennifer Lackey, is that the joint commitment account fails as an account of group *belief* since it cannot distinguish group *beliefs* from group *lies* and *bullshit*. The second is that the joint commitment account fails because it makes group belief *voluntary*, whereas genuine belief is *involuntary*. I propose an amendment to the basic joint commitment account which offers a unified reply to both objections. Although my novel account of group belief departs from the basic joint commitment account, it retains its spirit. The account entails that genuine group belief is much rarer than proponents of the joint commitment account have hitherto realized.

Keywords Group belief · Joint acceptance · Joint commitment · Aim of belief · Jennifer Lackey · Doxastic involuntarism

1 Introduction

Group belief ascriptions are ubiquitous. We use them to characterize group states and to explain group behavior. ‘Physicists believe they’ve detected the Higgs boson’, ‘The board of directors didn’t believe Ann made a compelling case for downsizing’, ‘The Catholic church sanctified John Paul II because it believes miracles occurred on account of his intercession’, ‘The US government invaded Iraq because it believed there were weapons of mass destruction’.

✉ Jeroen de Ridder Ph.D.
g.j.de.ridder@vu.nl

¹ Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

Understanding group belief is important, first of all, to make sense of such ascriptions: Should we take them literally or are they metaphors? It also matters for explanatory purposes. Just as we explain people's behavior with reference to their beliefs, so we explain the actions of groups with reference to their beliefs – both in everyday life and in social science. If there are no group beliefs, then explanations couched in terms of them are defective. Third, parity considerations also motivate inquiry into the nature of group belief. If the familiar project of understanding individual belief in epistemology and philosophy of mind is worthwhile, then, given that talk of group belief is just as normal and natural, the project of understanding group belief merits serious attention too.

One way to think of group beliefs is as shorthand for all or most group members having the same belief. Various philosophers have argued that this will not do, because sometimes a group believes something that few or even none of its members do.¹ They propose that there are robustly collective beliefs, i.e., beliefs that are genuinely the *group's* beliefs, rather than some mere combination of the group's members' beliefs.

I side with these philosophers here and propose and defend an account of robust group belief. My account emerges in response to two objections against an influential account of group belief, the *joint commitment account*. The next section introduces this account. In Sect. 3, I discuss Jennifer Lackey's objection that the joint commitment account cannot distinguish group beliefs from group lies and bullshit, and in Sect. 4 I show how it can be modified to deal with this objection. In Sect. 5, I introduce the objection that the account is unsatisfactory because it makes group belief voluntary rather than involuntary. Section 6 details that my modified account offers a straightforward response. Section 7 reflects on the limitations and scope of my novel account of group belief.

2 The joint Commitment account of group belief

Influential accounts of robust group belief are due to Margaret Gilbert (1987, 1989, 1994, 2004, 2014) and Raimo Tuomela (1992, 2004).² Their central idea is that groups form beliefs when their members jointly commit to letting a view stand as the group's view. Let's refer to this idea as the *joint commitment account* (JCA) of group belief. To have a specific proposal before us, here is Gilbert's most recent statement of her view:

The members of a population, P, collectively believe that p if and only if they are jointly committed to believe that p as a body. (...) The general idea of a joint commitment to @ "as a body" is this. The parties take it upon themselves to emulate, as far as is possible, a single body that @s, by virtue of the actions of each. (...) [To

¹ Gilbert (1989) and Tuomela (1992) are the seminal sources for this view; the next section gives more references.

² Other proponents are Schmitt (1994, 2014) and Rolin (2008) and critical sympathizers include Mathiesen (2006, 2011) and Wray (2007). Alternative accounts of robust group belief, which do not analyze it in terms of joint commitment, include Tollefsen (2003, 2015), List & Pettit (2011), and Pettit (2011).

form a joint commitment], each of the parties must express – in conditions of common knowledge – his or her readiness to be jointly committed with the others in the relevant way. (...) How are people to act so as to emulate, as far as possible, a body that believes that *p*? (...) They are to act as would any one of several *mouthpieces* of the body in question, thus uttering *its* beliefs, as opposed to the beliefs of any of its members, including the utterer. (Gilbert 2014, pp. 137–140)

I'll add a few clarifications and specifications. First, Gilbert's earlier work analyzed group belief in terms of joint commitments to *accept* a proposition.³ To accept a proposition is to 'go along with it', i.e., to choose a policy of positing or postulating it, to act as if it is true, and to use it as a premise in one's practical and theoretical reasoning (Cohen 1989, p. 368). The expressions 'emulating' and 'acting as mouthpieces' in the above formulation are functionally equivalent and nothings hangs on the terminology chosen according to Gilbert herself (2014, p. 140).

Second, as it stands, JCA is unable to handle groups with diverse roles in decision-making. Sometimes, designated members have special authority to determine what the group believes. For instance, a board of directors might determine what the company believes. To handle this, Tuomela (1992) introduces a distinction between *operative* and *non-operative* group members. On his account, group belief requires that operative group members satisfy conditions for group belief equivalent to Gilbert's, while non-operative members tacitly accept the group belief that is so formed (or at least ought to do so). I will take this modification for granted in what follows.

Third, on JCA, group belief is *irreducible* to the beliefs of individual group members (cf. Gilbert 1987; 2014, pp. 136–7). Shared individual belief that *p* among (operative) group members is neither sufficient nor necessary for group belief that *p*. For a group to believe that *p*, it is insufficient that most or all (operative) group members individually believe that *p* and that this is common knowledge among them. Something extra is required: The group has to make a decision to jointly commit to believe that *p*. Shared individual belief is also not necessary. A group can believe that *p*, even if no individual (operative) group member believes that *p*. Such scenarios have been called divergence cases, because the group belief can diverge from the beliefs of the group members. Several of the authors cited earlier offer the example of a job search committee. Every individual committee member might have a different belief about who the best candidate is, but share the belief that Alice is second best. In such a scenario, it is conceivable that they might jointly commit to the belief that Alice is the best candidate – and thereby come to believe it as a group.

Fourth, in order to jointly to commit to a view, a group must employ some procedure to decide what its view will be. This can be an officially established decision procedure such as voting, but also an informal or ad hoc deliberative process (cf. Gilbert 2006, 2014). The point is that the procedure generates the joint commitment and common knowledge required for group belief. All (operative) group members

³ Cf. Gilbert's original formulation: A group *G* believes that *p* if and only if the members of *G* jointly accept that *p*. Members of a group *G* jointly accept that *p* if and only if it is common knowledge in *G* that the individual members of *G* have openly expressed a conditional commitment jointly to accept that *p* together with the other members of *G*. (Gilbert 1989, p. 306)

must know that the other (operative) members agree to let a view stand as the group's view, conditional on similar commitments by the other (operative) group members.

Fifth, group members should individually endorse the procedure used by the group to form its belief.⁴ This need not require much and may even be implicit. Simply becoming a group member can be enough, when it is understood that membership involves endorsement of the designated procedure. If this condition is not satisfied, members may be unwilling to use the procedure and the conditions for forming a joint commitment will not be satisfied.

The core of view can be condensed into the following short statement, where 'joint commitment' and 'believing as a body' are to be understood as specified above:

(JCA) Group G believes that *p* if and only if G's operative members are jointly committed to believe that *p* as a body.

This will be the canonical view in what follows.

3 Group beliefs vs. group lies and bullshit

The first objection to JCA is due to Jennifer Lackey (2020).⁵ To appreciate her objection, we must recall that belief and truth are intimately connected. For you to believe that *p*, is for you to believe that *p* is true, to take *p* as true, or to regard *p* as being true. Unless you take *p* to be true, you don't believe it. This holds for false beliefs just as well. Even if *p* is in fact false, if you believe that *p*, you believe that *p* is true, albeit mistakenly.⁶ Joint commitment to believe, however, lacks such an intrinsic connection with truth. You can commit to emulating belief that *p* for many reasons other than your taking *p* to be true: because doing so seems convenient, because you want to explore the theoretical consequences of *p*, or because it serves some practical goal.

Lackey rightly demands of an account of group belief that it can distinguish group *beliefs* from group *lies* or *bullshit* (in Harry Frankfurt's sense⁷). JCA cannot do this and must therefore be rejected. Lackey presents the following examples of a group lie and of group bullshit to support her case.

Tobacco Company.

Phillip Morris, one of the largest tobacco companies in the world, is 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. While the members

⁴ Or, at the very least, a significant majority of them should.

⁵ Lackey targets older formulations of Gilbert's account, couched in terms of joint acceptances. I have rephrased her criticisms to apply to Gilbert's most recent version.

⁶ This idea has often glossed by saying that belief aims at truth. Thus Williams (1973, pp. 136–7): "When I say that beliefs aim at truth, I have particularly in mind three things. First: that truth and falsehood are a dimension of an assessment of beliefs as opposed to many other psychological states or dispositions. ... The second feature under this heading: to believe that *p* is to believe that *p* is true. ... The third point ... is: to say 'I believe that *p*' itself carries, in general, a claim that *p* is true." Cf. also Velleman (2000).

⁷ Frankfurt's official characterization of a *bullshitter* is that he "is neither on the side of the true nor on the side of the false. His eye is not on the facts at all, as the eyes of the honest man and the liar are, except insofar as they may be pertinent to his interest in getting away with what he says. He does not care whether the things he says describe reality correctly. He just picks them out, or makes them up, to suit his purpose" (Frankfurt 2005, pp. 55–56).

of the board of directors of the company believe this conclusion, they all jointly agree that, because of what is at stake financially, the official position of Phillip Morris is that smoking is neither highly addictive nor detrimental to one's health, which is then published in all of their advertising materials. (ibid., p. 195)

Oil Company.

After the oil spill in the Gulf of Mexico, BP began spraying dispersants in the clean-up process that have been widely criticized by environmental groups for their level of toxicity. In response to this outcry, the executive management team of BP convened and its members jointly accepted that the dispersants they are using are safe and pose no threat to the environment, a view that they then made public through all the major media outlets. It turns out that BP's executive management team arrived at this view with an utter disregard for the truth – it simply served their purpose of financial and reputational preservation. (ibid., pp. 197–8)

In both examples, the group plausibly satisfies JCA's conditions for group belief. The operative group members jointly commit to emulating, as far as possible, a single body that believes a position. Group members have openly expressed their willingness to let this position stand as their official position, so that this is common knowledge among them; they have used the designated deliberative procedure to form their position; they act as if it's true, etc. Hence, JCA rules that Philip Morris *believes* that smoking is neither addictive nor unhealthy, and that BP *believes* that their dispersants are safe. This seems wrong.

Lackey considers and rejects two possible replies (ibid., pp. 200–1). First, a proponent of JCA may urge that *Tobacco Company* should be construed as a case in which a group *first* jointly commits to believe a proposition *p* and *then* also jointly commits to spread it about that not-*p*. This makes the case relevantly different from those in which a group straightforwardly commits to believe that *p* without qualification. Hence, it is possible to distinguish between group belief and group lies after all.

Lackey rightly finds this reply wanting. It may be possible to interpret the case in this manner, but this is neither necessary nor natural. A group might as well commit to emulate belief that not-*p* directly, without the intermediate step of first committing to *p*. The fact that all group members in fact believe that *p* doesn't show that such an intermediate step is necessary, since, as we saw above, universally shared belief is not sufficient for group belief on JCA. Furthermore, in cases of group bullshit, there simply cannot be any preceding commitment to a proposition that would constitute the 'real' group belief. For it is characteristic of bullshit that the subject lacks prior belief either way. Hence, there is nothing preceding the joint commitment to emulate belief that could establish a relevant difference with group belief.

A second possible reply is to grant that JCA is incomplete as it stands, but to argue that it can easily be repaired. In the counterexamples, a group jointly commits to acting as if they believe a proposition, but it does so for reasons that have nothing to do with truth. This shows that not just any joint commitment to emulate belief leads to group belief. A further condition is needed: only joint commitments to emulate belief *that have a proper connection to truth* generate group belief.

In response, Lackey contends that no truth-connection condition can do the job. She gives two reasons. First, she notes that allegedly paradigmatic cases of group belief from the literature in fact lack a proper connection to truth. The job search

committee that accepts the candidate whom all committee members believed to be second best, is more likely to have been motivated by the need to reach consensus or to avoid collegial conflict, than by a concern to find out which candidate is *truly* best. Something similar holds for another oft-cited case, in which an association issues a formal statement. The American Academy of Pediatrics jointly commits to the view that there is no causal connection between vaccination and autism and communicates this view in all its materials. We can imagine that they might do this even when there is uncertainty about the evidence and ongoing controversy among members. Lackey avers that the medical association is motivated by a concern for public health, not for the truth. Hence, this allegedly paradigmatic case, too, fails to satisfy a condition linking JCA-style group belief to truth.

Lackey's second reason for thinking that no suitable truth-connection condition exists, is that various processes that lack "sensitivity to truth, or to the available evidence, or to some other epistemically proper feature" (ibid., p. 200) can produce beliefs, both at the group and the individual level. Wishful thinking, for example, is not sensitive to truth, but does produce belief. The same can be said about a group-level process like group polarization (Sunstein 2002). Since an account of group belief ought to leave room for beliefs that are produced by processes that lack a connection with some epistemically proper feature, any truth-connection condition that is added to JCA will rule out too much.

4 Response: A better truth connection

I will now argue that Lackey's objection to JCA fails – or at least that a modified version of JCA that is independently motivated avoids her objection. Key to my response is the addition of a truth-connection condition that mirrors more closely how individual belief is connected to truth. Recall that, in support of her claim that no suitable truth-connection condition can be found, Lackey noted that various processes that lack "sensitivity to the truth, or to the available evidence, or to some other epistemically proper feature" (2020, p. 200) can produce beliefs. Hence, adding a condition to JCA stipulating that the process for forming a group belief needs to be 'sensitive to truth, to the available evidence, or to some other epistemically proper feature' would inappropriately rule out legitimate belief-forming processes.

But Lackey's quick suggestions are nonstarters, because they don't apply to individual belief either. Individuals often form false beliefs, so individual belief isn't straightforwardly sensitive to the truth. Lackey herself already noted that individuals sometimes form beliefs through wishful thinking, cognitive bias, or for no clear reason whatsoever, so individual belief isn't always sensitive to the available evidence either. Rather, belief requires nothing more or less than that a proposition *strikes you as true*, that you *regard it to be true*, or that it *seems true to you*. Of course, what strikes you as true, what you regard as true and what seems true to you is not always sensitive to truth, the available evidence, or some other epistemically proper feature. A plausible condition on group belief must capture the analogue of a proposition's striking someone as true. A straightforward way of cashing this out would be to require that the accepted view strike all (or most of) the operative group members as

true, or simply that they all believe it to be true. Thus, one truth-connection condition (TC) could be:

(TC1) Group G believes that p only if p strikes all the operative members of G individually as true, or, alternatively, only if all of the operative members of G individually believe that p.⁸

The full modified JCA would thus become:

(JCA-TC1) Group G believes that p if and only if G's operative members are jointly committed to believe that p as a body, and all of the operative members of G individually believe that p.

Because joint acceptances that satisfy TC1 parallel individual belief closely, they plausibly count as instances of group belief. Note that the addition of TC1 immediately gets rid of Lackey's counterexamples, for both in *Tobacco Company* and *Oil Company*, it is not the case that the view that is accepted strikes all operative group members as true. The cases are set up so that the operative group members individually disbelieve the accepted view or lack belief on it either way.

Even though JCA-TC1 handles Lackey's objection, proponents of JCA will be reluctant to embrace it. It has the unhappy consequence of making universally shared individual belief a necessary condition on group belief, while this was something they explicitly sought to avoid. A condition that does not require this is preferable.

Since, on JCA, group beliefs involve a procedure for generating joint commitment, we can look for a condition that involves this procedure, rather than the view resulting from it. Here is a proposal:

(TC2) Group G believes that p only if G's operative members individually believe that the procedure they use to form the joint commitment to believe that p as a body is reliable, i.e., likely to lead to true outputs.⁹

Adding this to JCA, we get:

⁸ Perhaps we need not require that *all* operative members believe that p. Especially if the number of operative group members is large, a designated majority might suffice. I will omit this qualification in what follows.

⁹ TC2 bears some resemblance to a condition that Kay Mathiesen formulates for a group to be 'subjectively epistemically rational': "A group G is subjectively epistemically rational iff G has the goal of believing* truths and avoiding falsehoods and *abides by epistemic practices that G believes* are effective in achieving that goal*" (2006, pp. 165–6, my italics), where 'belief*' is defined as 'having the view that'. However, there are at least three problems with Mathiesen's condition as it stands. (1) She leaves it open whether being subjectively epistemically rational is a necessary condition for group *belief* or for truth-aimed group *acceptances*. Hence, the point at stake in Lackey's objection remains unresolved. Taken as a condition on group belief, it faces two further problems. (2) By folding a truth-connection condition into a subjective epistemic rationality condition, her account rules out subjectively irrational beliefs by stipulation. This is undesirable: Epistemic rationality and the connection between belief and truth are separate issues. It may be controversial whether epistemic akrasia exists, but it should not simply be defined out of existence (cf. Greco 2014 and Horowitz 2014 for discussion). (3) Most importantly, however, the condition is circular because it employs the very notion of belief (or 'having the view') it is supposed to elucidate. By making it a condition on group belief that a group *believes* (or has the view) that its epistemic practices are effective in reaching the epistemic goal, Mathiesen's condition fails to give us a handle on what it is for a group to believe something in the first place. (Mathiesen also discusses a structurally parallel condition for 'objective epistemic rationality', where the second clause includes the stipulation that the group's epistemic practices *are* in fact effective (2006, p. 166). This clearly isn't a plausible condition on group belief. If the way the group forms its beliefs is in fact reliable, but group members don't believe this to be so, the resultant group state lacks any connection with seeming truth.)

(JCA-TC2) Group G believes that p if and only if G's operative members are jointly committed to believe that p as a body, and individually believe that the procedure they use to form the joint commitment to believe that p as a body is reliable, i.e., likely to lead to true outputs.

Unlike TC1, TC2 does not require universally shared belief that p among group members, but still secures the desired result that not any joint commitment to emulate belief gives rise to a genuine group belief, which was the source of trouble in the basic account.

TC2 must be clarified further.¹⁰ First, talk of reliable procedures raises a version of the familiar generality problem for reliabilism (Conee & Feldman 1998). Whether or not a procedure is reliable depends on what that procedure is, or, in other words, how it is supposed to be individuated. For instance, is the procedure 'taking a vote', 'the current operative group members taking a vote', or 'the current operative group members taking a vote about topic T', etc.? Reliability will vary depending on what option we choose. The generality problem is a problem for any view that involves a clause about the reliability of belief-forming processes at either the individual or collective level, so a full discussion is beyond the scope of the present paper. However, Jeffrey Dunn (2021, pp. 5661–2) has recently made a case that the following items must factor into a plausible individuation of group procedures: (i) group constitution, (ii) method, and (iii) topic. This is to say that (i) changes in (operative) members result in a different procedure; (ii) changes in the methods used for forming joint commitments (voting, deliberation, etc.) result in a different procedure; and (iii) a change in topics about which a view is formed results in a different procedure. This is motivated by the observations that (i) the identities and number of group members affect the reliability of a group procedure; (ii) the nature of commitment-forming methods affects their reliability; and (iii) individuals and methods aren't equally reliable for any and all topics.¹¹

Second, TC2 stipulates that G's operative members must believe that the procedure they use is reliable, where this procedure is the one individuated in the sense just described. But this doesn't mean that all operative group members must have consciously reflected on the truth-conduciveness of their procedure (although that is also possible); it suffices if they would answer affirmatively, were they asked about its reliability, and if they lack reason to think that the procedure is unreliable. The condition thus even allows that individual group members can have different conceptions of how the procedure is to be individuated, as long as they would assent to the claim that the procedure as individuated along the above lines is reliable.

Third, the reliability in question must be understood in a *modal* sense rather than a track-record sense. Group members need not believe that the procedure they use has led to a high ratio of true over false outputs in the past – after all, some procedures might have a very limited track-record. It is sufficient that they believe that it is likely

¹⁰ Thanks to Jennifer Lackey and two anonymous referees for pushing me to think through these points.

¹¹ Some vagueness admittedly remains. For example, is deliberating for 30 min the same procedure as doing it for 35 min – and if it is, after how much time does one draw a line? Is simple majority voting the same procedure as voting which requires a 55% majority? But, I submit, such vagueness is tolerable since we don't need – and can never give – exact reliability estimates anyway.

to produce a correct on this occasion in the actual world and a broad range of nearby possible worlds.¹²

Fourth, we can distinguish between a narrow and a broad sense of the reliability of a procedure. In the narrow sense, it only concerns the process by which inputs are transformed into outputs. For instance, the process of casting and counting votes in a simple majoritarian aggregation procedure.¹³ In the broad sense, reliability is concerned with the quality of the inputs as well as the quality of the way in which they are processed.¹⁴ TC2 is concerned with reliability in the latter, broad, sense. To see why, consider that the (operative) group members may well believe of a specific procedure that it is reliable in the narrow sense, but also believe that, on this particular occasion, it fails to be reliable in the broad sense because its inputs happen to be mostly false or unreliably formed beliefs. (Perhaps the group is considering taking a vote after a bibulous dinner.) In such a case, the operative group members all have reason to believe the accepted view was unreliably formed and likely to be false. This precludes the view from being a genuine group belief, just as in the individual case. An individual who is aware that her view was formed by an unreliable process will not take that view to be true, or believe it. If a procedure is only believed to be reliable in the narrow sense, then, the truth-connection isn't strong enough.

For an example of JCA-TC2 in action, consider the following case:

Philosophy Department.

The members of the Philosophy Department are figuring out whether becoming part of a broader Humanities School will make them more likely to survive budget cuts. To reach a decision, they have agreed to share and discuss their individual takes on the issue, correct any misunderstandings, draw up a big table to weigh and score the pros and cons, and then determine a final assessment of the two options. The option that comes out of this procedure as most plausible is then accepted as the Department's view. All faculty members believe that this deliberative procedure is reliable, because their individual perspectives on the topic are, although well-informed, also limited and they take thorough deliberation to be a reliable method for pooling their information and aggregating it appropriately into an assessment.

In addition to satisfying the conditions of the basic account, this scenario also satisfies TC2. The Department thus comes to *believe* the outcome of its procedure.

Next, note that TC2 gives the right verdict in *Tobacco Company* and *Oil Company*. In both cases, the group members did not believe that the procedure they used to form the group's position was an epistemically reliable one. Group members were quite aware that it was unreliable; that was the very point. Hence, no group belief is formed in these cases. JCA-TC2 distinguishes successfully between group beliefs on the one hand and group lies or bullshit on the other.

Let's consider how JCA-TC2 fares on closer inspection. Do Lackey's two reasons to think that no suitable truth-connection condition can be given spell trouble for

¹² Note that known track-record reliability of a procedure might nonetheless be a good reason to believe the procedure to be reliable in the modal sense.

¹³ We can allow that narrow reliability may be undefined for various kinds of procedures, since everything may depend on the epistemic quality of the inputs.

¹⁴ See Goldman 2014 and Dunn 2021 for a similar distinction.

TC2? The first was that the (allegedly paradigmatic) examples of group belief offered by the proponents of JCA fail to display the right sort of connection to truth. But this is overhasty. With the above candidate truth-connection condition (TC2) in mind, we see that cases like those of the job search committee or the medical association are crucially underspecified, for they don't make it clear what the committee members or the medical doctors think about the reliability of the procedure they use to form their respective views. If the members of the job search committee believe that casting votes is an epistemically reliable procedure to reach a conclusion about who the best candidate is, then their joint commitment satisfies TC2 and the resulting group view is an instance of group belief. If they don't believe this and only take a vote because, say, it's in the department's rules or the only way to avoid collegial conflict, then the resulting view is not a group belief. The same goes for the medical association. If the operative members of the Academy think that their deliberation is an epistemically reliable way of forming a position, then their official position is an instance of group belief.¹⁵ If not, then not.

What's more, even if these cases ultimately fail to constitute plausible examples of group belief, nothing serious follows with regard to the existence of robust group belief. The only upshot would be that proponents of JCA have been offering poorly specified examples to motivate their view. That's embarrassing, but it doesn't mean that there are no good examples.

Lackey's second reason for thinking that no suitable truth-connection condition exists was that any such connection was going to rule out too much, because cognitive processes that are insensitive to the truth or another epistemically proper feature can also lead to beliefs. Again, TC2 can accommodate this. It only requires that operative group members *believe* that the procedure used in jointly committing to a view is reliable. Their beliefs can be false. Groups can fall prey to wishful thinking, bias, polarization, and other distorting influences without their members realizing it. In such cases, the operative group members may believe that they are using a reliable procedure even when they in fact aren't. Thus, TC2 is satisfied and the resulting joint commitment constitutes a group belief.

Let's consider two new objections to my proposal. First, a purported counterexample to JCA-TC2, a variation on *Tobacco Company*:¹⁶

Tobacco Company 2.

Each individual operative member of a tobacco company board has excellent evidence that smoking increases risk of heart disease. But they want to hide this. They locate a reliable and well-respected academic who, in addition to having published a large body of high-quality reliable work, has cast doubt on the link between smoking and heart disease in his official publications and public performances. The members opt to form the company's beliefs by deference to this academic. Each operative member believes that this is a reliable procedure, given the academic's general reliability. The board thus comes to believe that it is doubtful that smoking increases the

¹⁵ Note that their believing this doesn't exclude their also being motivated by a concern for public health. Presumably, finding out the *truth* about side-effects of vaccines is instrumentally valuable for public health.

¹⁶ Thanks to an anonymous referee for coming up with the example.

risk of heart disease despite each individual board member having strong evidence of the increased risk.

The objection is that JCA-TC2 gives the wrong result, just like JCA did for the original *Tobacco Company* case. The board satisfies the required conditions, so JCA-TC2 rules that the board believes there is doubt about the connection between smoking and heart disease. This is incorrect, because the board is clearly (and intentionally) lying.

This is too quick, however. Like some of the earlier cases we encountered, *Tobacco Company 2* is ambiguous. To establish whether board members indeed satisfy TC2, more must be said about how they weigh the strength of their own evidence vs. that of the academic's expert testimony, because that determines how they assess the modal reliability of the chosen procedure on this particular occasion. There are two options. The first is that they take their own excellent evidence to be compelling and the expert's testimony on this topic misleading. (Note that this is perfectly compatible with believing the academic to be generally reliable and reliable on other topics and occasions.) The second option is that they think the expert's testimony trumps their own evidence on this occasion. (They might assume that the expert is surely aware of their evidence but also has further countervailing evidence which makes him believe there is reason for doubt). On the first option, board members don't satisfy TC2, because they don't take the procedure of deferring to this expert's testimony to be modally reliable on this occasion. Hence, JCA-TC2 rules that the board forms no genuine group belief (even though they may commit to emulating belief in their actions). On the second option, board members do satisfy TC2 and hence the board forms a genuine belief. This is so in spite of their original intention to want to hide the evidence and lie. In this scenario, most or all board members would probably even give up their individual belief about the risks of smoking for heart disease because they are under rational pressure to accept expert testimony. In spite of initial appearances, then, *Tobacco Company 2* doesn't undermine JCA-TC2.

Here's a second worry for JCA-TC2. Earlier, I rejected TC1 because it re-introduced universally shared individual belief as a necessary condition on group belief. One might object that TC2 implicitly does the same thing and that, therefore, friends of JCA will not find it palatable either. If all the operative group members believe that the procedure they use to form a joint commitment is reliable, then individual group members come under rational pressure to adopt the group belief as their individual belief and relinquish any dissenting beliefs they may have. Compare: if I believe my friend to be a reliable informant about a topic, then, absent any special reasons to think she's mistaken on this particular occasion, I am under rational pressure to revise my beliefs about that topic in accordance with her sincere testimony. So even though TC2 doesn't make universally shared belief an explicit condition on group belief, it has the same effect.

This objection fails. Contrary to what it suggests, TC2 does leave some room for individual dissent, although a joint commitment satisfying TC2 indeed creates rational pressure on dissenters to conform their beliefs to the group belief. The pressure, however, can sometimes be resisted, although this is admittedly a fine line. A dissenter might believe that, even though the procedure used by the group was *likely* to produce a correct output, it didn't do so on this occasion by some stroke of bad luck.

She might feel that, even though she shared her views with the other group members, they inadvertently weren't fully grasped or accidentally weren't given their proper weight in reaching the final decision. (But note that this cannot be because there is some systematic problem with the procedure used. For that would prevent the dissenter from believing the procedure to be reliable in the modal sense on this occasion, in which case TC2 wouldn't be satisfied.) Alternatively, she may simply find herself psychologically unable to believe what the group believes. All of this is compatible with her continuing to regard the procedure as reliable.¹⁷ The situation parallels that of a layperson who receives expert testimony. In most cases, the layperson should believe what the expert tells her, but sometimes she can demur. She may have reason to think that the expert's testimony inadvertently misfired on this occasion, even though it was likely to be correct. Or she may be psychologically unable to believe what the expert tells her. This is compatible with her continuing to regard the expert as reliable on this occasion.

These reflections do show that so-called divergence cases, where a group believes something that not all (or even none) of its members do, may be rarer than defenders of JCA thought. The requirement that operative group members must believe the commitment-forming procedure to be reliable (as per TC2) generates rational pressure on group members to conform their individual beliefs to the group's belief, which can only be withstood in fairly exceptional circumstances. Are we then back to (almost) universally shared belief as a necessary condition on group belief after all? Not quite. An important difference remains between JCA-TC2 and summative analyses of group belief, which is obscured by the language of necessary conditions. The difference is that the order of explanation of the two accounts is exactly opposite. On summative analyses of group belief, a group believes that *p* *because* most or all of its (operative) members believe that *p*. On JCA-TC2, in contrast, all or most of the operative group members will come to believe that *p* (if they didn't already do so), *because* the group has formed the belief that *p* by means of a procedure that the group members believe to be reliable.

Summing up: there is indeed a difference between individual belief on the one hand and joint commitment to believe as a body in the sense of the basic JCA on the other. The former has an intrinsic connection with seeming truth that the latter lacks. But Lackey's conclusion that no analysis of group belief along the lines of JCA can succeed is too quick. JCA can be fixed by supplementing it with a truth-connection condition, such as TC2. JCA-TC2 holds that only a subset of joint commitments to believe as a body result in genuine group beliefs, namely those that satisfy TC2.

¹⁷ Mathiesen (2011) gives an argument for a similar conclusion, to wit that a group's belief can sometimes diverge from the shared beliefs of the group members and that we therefore ought to treat groups as epistemic agents in their own right. Her argument, however, is entirely case-based and doesn't develop a systematic account of group belief. Moreover, her key piece of evidence is a case which involves a scenario where a group's belief and its members' beliefs diverge only for a short period, namely the time between (a) the moment at which group members have finished giving their input to the procedure that's used to determine the group belief and (b) the moment when the results of this procedure become known to the group. Even if we grant – which I'm not sure we should – that this is a plausible case of a group belief diverging from group members' individual beliefs, it certainly is a rare and idiosyncratic one. If robust group belief were only instantiated in such cases, it would be highly elusive and of little independent interest.

Lackey's objection thus teaches us an important lesson. It's not that JCA fails as an analysis of group belief, but that genuine group beliefs are rarer than the friends of JCA have hitherto assumed.

5 Voluntary control

The second objection pushes another way in which joint commitments to believe may be too dissimilar to individual belief to make JCA a plausible account of group belief. The worry is that JCA implies that group belief is under direct voluntary control, whereas it is a widely held view among epistemologists that this is not the case (cf. Alston 1988; Feldman 2001). Hence, analyzing group *belief* as a form of joint commitment is problematic.¹⁸

This objection has been put forward in Wray (2001) and Meijers (2002), but Raul Hakli (2006) develops it in most detail. He analyzes different types of collective propositional attitudes¹⁹ and distinguishes different senses of voluntary control. The notion at stake in the debate about the voluntariness of belief is that of *direct voluntary control*. Having voluntary control over belief is having an ability to believe, disbelieve, or suspend judgment *on a specific proposition* and to do so *at will*. Although we sometimes have *indirect* voluntary control over what we believe (I can 'choose to believe' that the door is open by opening it and observing the result of my action), we cannot *directly* choose to believe, disbelieve, or suspend judgment on a proposition of our choosing. Furthermore, although we do have some direct voluntary control over *which propositions we consider* and over *whether we gather evidence on a specific proposition and, if so, how much*, we do not thereby have direct control over our beliefs. What we end up believing when we consider a proposition or gather evidence is determined by how things seem to us upon doing so, and that is not something we choose. Hakli summarizes the upshot of his discussion as follows:

Doxastic states [i.e., propositional attitudes, JdR] can be classified into voluntary and involuntary. Involuntary doxastic states can be called proper beliefs and voluntary ones can be called acceptances. [...] The phenomenon that has been called collective belief or group belief, when understood in a non-summative fashion [i.e., along the lines of JCA, JdR], falls into the category of voluntary doxastic states. (Hakli 2006, pp. 294, 297)

He concludes that JCA is misguided. Joint commitments to believe as a body are voluntary and therefore it is a mistake to speak of group *belief* when it is construed along the lines of JCA.

Before we evaluate the objection, it is instructive to consider more carefully how joint commitments to believe are supposed to be voluntary. Clearly, they are not

¹⁸ Note that, even if one does think that belief is under significant control, this objection still needs a reply. First, if JCA can only succeed if it relies on the controversial view that belief is under voluntary control, that makes it significantly weaker than it would otherwise be. Second, if belief is under voluntary control, the objection can be recast as insisting that JCA implies that group belief is under *more* voluntary control than belief and thus it makes group belief too dissimilar to individual belief to be plausible.

¹⁹ Hakli himself writes about *doxastic* states. But since that term is semantically connected to *belief*-like attitudes, I prefer the more neutral '*propositional attitudes*'.

under any individual group member's direct control.²⁰ But they are arguably under the group's direct voluntary control. If we think of the group's will as constituted by its collectively endorsed use of the designated decision procedure, it's clear that a group can come up with a proposition *p* and decide to commit to emulating belief that *p* by going through its procedure.²¹ This does not mean that a group can directly commit to any proposition whatsoever. Perhaps some propositions are such that a particular group would never commit to them in virtue of its deliberative dispositions or voting preferences.²² The point is that *some* propositions are such that the group can directly decide to commit to them. That suffices to show that group belief, construed as joint commitment, is sometimes under direct voluntary control. Since, according to the objection under consideration, genuine belief is not (or at least less so, cf. note 18 above), JCA is unacceptable as an account of group *belief*.

We can bring the objection into sharper relief by comparing actions with beliefs and commitments respectively. Choosing actions is characteristically voluntary at two levels. First, you can choose your goals in acting: what do you want to achieve? Second, you often have a choice among a range of possible actions to achieve your chosen goal. Not every action is preceded by conscious choices of both these types. Some goals may be inevitable, but many aren't. And sometimes there really is just one feasible action to achieve a chosen goal, but often multiple actions are available. So, for at least some of your actions, you have direct voluntary control at two levels: that of choosing the goal and that of choosing the action.

Many philosophers maintain that beliefs are different. As we noted in Sect. 3, belief is intrinsically connected to truth (e.g., Velleman 2000; Wedgwood 2002; Shah & Velleman 2005).²³ Hence, truth is not an optional goal for beliefs; you aren't free to decide the goal for which you believe things. You cannot choose to believe something because, say, it makes you happy. So there is no voluntary control at the level of the goal of our beliefs. At the level of adopting specific beliefs, there is some indirect voluntary control, however, because you can choose which propositions to consider. With beliefs, then, there is voluntary control at at most one level.

In this respect, commitments are more like actions than beliefs. You can choose both the goal for which you commit to a proposition – in so far as it is a goal that can

²⁰ Barring dictatorial decision procedures in which one operative group member calls the shots. As this is a limiting case of group belief, I will henceforth ignore this qualification.

²¹ Gilbert's own reply to the voluntariness objection (Gilbert 2014, pp. 153ff) is to deny this. According to her, collective belief is not under *direct* voluntary control, because it necessarily involves a two-step procedure in which group members must *first* collectively form a desire or intention to believe a proposition and *then* jointly commit to emulate belief in that proposition. This is correct as far as it goes, but it doesn't get to the heart of the worry. The mere fact that the process has two steps doesn't show that it is not under direct voluntary control. The point remains that a group can choose what to believe entirely on its own accord, albeit through a two-step process. See also Gilbert & Pilchman (2014, pp. 205ff) for a follow-up response to Hakli (2006).

²² The same holds at the individual level. Not anyone can accept any proposition whatsoever. Some propositions may be beyond an individual's grasp or too bizarre or repugnant to accept.

²³ Nothing in what follows will depend on this particular formulation of the connection between belief and truth. If, as some have argued (e.g., Williamson 2000; Adler 2002), *knowledge* rather than truth is the norm for belief, then that's fine too. Since knowledge is factive, the required connection with truth remains in place.

be served by committing to a proposition – and the specific proposition you commit to. For example, you can commit to a proposition because it suits your practical interests (as in Lackey’s *Tobacco Company* and *Oil Company* cases), because you want to explore its theoretical or practical ramifications, or because it seems true to you. Once you have chosen a goal, you also have voluntary control over which propositions to consider for committing to. So, like actions but unlike beliefs, commitments are voluntary at two levels: that of their goals and that of their specific content. Therefore, says the objector, JCA’s construal of group belief fails.

6 Response: JCA-TC2 makes group beliefs involuntary

I will now show that my modified account of group belief offers a straightforward reply to the objection from voluntariness. Before I do so, however, I briefly consider and discard another reply, which may seem promising.

One could suggest that the difference in voluntariness between belief and joint commitment to believe is merely superficial and doesn’t amount to a decisive blow against JCA. The involuntariness of belief could just be a contingent feature of belief *as it occurs in individual humans* and not essential to belief as such. Perhaps a different psychological make-up would have given us direct voluntary control over our beliefs (cf. Alston 1988, p. 263). Or perhaps *human* belief is necessarily involuntary, but not belief in general. Aliens or God might be able to believe things at will or, more to the present point, maybe groups can (cf. Tollefsen 2002, pp. 96–98; Gilbert & Pilchman 2014, pp. 207ff).

I don’t think this is compelling. First, it’s too speculative. We have little or no evidence about what sorts of mental actions and attitudes a different psychological make-up would enable. In fact, it’s hard to see how we could acquire such evidence, let alone for what God or aliens can do.²⁴ Second, it appears ad hoc or even question-begging. The issue at stake is whether JCA forms a plausible account of group belief. The objection from voluntariness holds that it is not, because joint commitment is voluntary and belief is not. To then propose that involuntariness is merely a contingent feature of *individual* belief, without supplementing this proposal with a principled explanation for why group belief would lack this feature, is ducking the issue. The objector is owed at least a sketch of an account of belief that can explain why individual belief is involuntary and group belief not.²⁵

²⁴ Furthermore, Bernard Williams (1973) has argued that belief is *necessarily* involuntary.

²⁵ Yet another reply would be to argue that, contrary to popular philosophical opinion, we do have direct voluntary control over (a) some (cf. Ginet 2001; Nickel 2010) or (b) all of our beliefs (cf. Montmarquet 1986; Steup 2012). Full treatment of this idea would require a lengthy excursion into the literature on doxastic involuntarism. Let me just point out two things in response. First, even if it were true that some beliefs are under direct voluntary control (cf. McHugh 2012; Peels 2014 for arguments against this), this does little to undermine the general dissimilarity between beliefs and joint commitments. After all, the worry was that joint commitments to emulate belief are *always* voluntary and therefore too different from beliefs. Showing that there is a small subset of voluntary beliefs doesn’t alleviate this worry. Second, establishing that all our beliefs are under voluntary control is an extremely tall order. Both everyday experience and received philosophical wisdom speak against this. (Peels (2014) rebuts Steup’s (2012) case for this idea.)

Let's turn to a better response. We can concede the objector's points that joint commitments to believe as a body resemble voluntary actions more than involuntary beliefs and that this difference matters. However, not *all* joint commitments to believe are voluntary. The core of my reply to the objection from voluntariness is that joint commitments which also satisfy TC2 are effectively involuntary in much the same way as individual belief. Hence, *those joint commitments that qualify as group beliefs are not under direct voluntary control.*

Recall that TC2 contains a clause about individual operative group members' beliefs. They must believe that the procedure used to arrive at a group view is reliable. Individual beliefs are not under anyone's direct voluntary control – this is the foundation of the objection. Satisfying TC2 is thus not under the group's (or the group members') direct voluntary control and therefore the joint commitment as a whole isn't either. So the objection from voluntariness fails as an objection against JCA-TC2.

Let's consider potential objections. First, whether or not they believe a procedure to be reliable is not under the group members' direct voluntary control, but the choice for a particular procedure is. Since groups can directly choose their commitment-forming procedure, they can sometimes foresee that a procedure that seems reliable to them will lead them to jointly commit to a certain view. For instance, if they can choose to adopt majority voting as their procedure, think that this is reliable, and furthermore know how the majority will vote, they could, in effect choose to adopt a certain group belief by choosing to adopt a voting procedure. So even on JCA-TC2, some group beliefs are under voluntary control and therefore this account too succumbs to the voluntariness objection.

This objection fails, because the sort of control described here is *indirect*. The group doesn't choose to adopt a specific belief directly, but it chooses to do something else – to wit, adopt a procedure – which it knows will lead to the adoption of a specific belief. We have this sort of indirect control over individual beliefs, too. For instance, when we have direct control over our evidential situation. We can 'choose to believe' that the light is on by turning it on.

It may be objected further that, although these are both instances of indirect control, there is a difference. The group does not influence its evidential situation in order to form its belief. Instead, it uses foreknowledge that a procedure will lead to a given belief. I agree that this is a difference, but I don't see why it would matter. Since individual belief formation does not involve something like the use of a procedure, it's impossible to give an example of an exactly parallel phenomenon for individual belief. But the following should come close enough to see that there is nothing to worry about here. Sometimes, we know in advance that performing certain actions – reading a book or talking to people – will lead us to adopt a specific belief. For instance, suppose I have observed that friends with moral sensibilities similar to mine all became vegetarians upon reading Jonathan Safran Foer's *Eating Animals*. I haven't read it yet. Nonetheless, I know that, once I get around to it, I will form the belief that eating animals is morally unacceptable.²⁶ Hence, there is a sense in which I

²⁶ For those inclined to deny that I have genuine foreknowledge here, you can make the evidence as strong as you like. Perhaps there's a strong correlation between my past belief changes and those of my friends,

can decide to adopt that belief, namely by deciding to read the book. But that doesn't mean I have direct voluntary control over my belief. Once I have read the book, it will seem to me that eating animals is wrong, so I will believe it. But at that point, I cannot simply decide to believe otherwise, so I don't have *direct* voluntary control over my specific belief. The same applies to the group. Once it has gone through the procedure, its belief becomes fixed and it cannot *then* decide to believe something else. But this is really to say that, in the envisioned scenario, the group's belief isn't under its *direct* voluntary control.

Moreover, in the envisioned scenario any voluntary control that the group has is contingent on two factors outside the control of the group: first, that group members believe some procedure to be reliable and, second, that the voting behavior of a majority of group members happens to be known in advance. The group can neither make its members' beliefs about the reliability of a procedure line up in the right way, nor can it determine how its members will vote, since that, too, depends on their individual beliefs.²⁷ So not only is the control that the group can exercise over its belief indirect, it also depends on contingencies outside the group's control. We could compare the situation to that of someone who happens to learn that her brain is wired so that if she were to perform some action, she would end up believing proposition *p*. This person can voluntarily choose to believe *p* (namely, by choosing to perform the action), but since her control is both indirect and dependent on a contingency about her brain, which she just happens to know about, the case fails to constitute a counterexample against doxastic involuntarism. Similarly for group belief: cases like the one envisioned do not threaten the involuntariness of group beliefs as construed by JCA-TC2.

Even if all of this is granted, there is a further difference in voluntariness between individual belief and group belief, which may be thought to undermine the involuntariness of JCA-TC2-style group belief.²⁸ Whereas *refraining from belief* is not under direct voluntary control for individuals – they form beliefs automatically and cannot help doing so – it is for groups. More precisely, while individuals can voluntarily choose whether or not to consider the evidence for a proposition, once they do so, they cannot voluntarily refrain from belief, disbelief, or suspension of judgment anymore. This is different for groups: even when group members have considered the evidence, they can voluntarily choose not to form any group belief (or other state) by not employing their commitment-forming procedure.

In response, first note that this objection changes the topic. The original involuntariness objection held that, because individual belief is involuntary, an account of group belief should make group belief involuntary too. JCA-TC2 does this; a group cannot decide at will what to believe. But now, the objection is that an account of group belief should also make *belief formation* involuntary. We have gone from the involuntariness of what to believe to that of whether to believe. This may or not be an

perhaps I'm always strongly inclined to believe the conclusions of the books I read, etc.

²⁷ That is, as long as the group members are to believe that the voting procedure is reliable. A group (or some of its members) can influence how members will vote (bribery, intimidation, etc.). But once it does so, group members will no longer believe it to be reliable and any resulting acceptance will fail TC2.

²⁸ Thanks to an anonymous referee for raising this objection.

important dissimilarity, it's not the one that has been at stake in the debate about the voluntariness of group beliefs and acceptances so far.

Of course, the mere fact that an objection is novel is not a good reason for dismissing it. So my second and more substantive reply questions the force of the objection directly. Why should the mere fact that a group could have refrained from forming a state that is purportedly a group belief preclude that state from being a group belief? Imagine that future neuroscience – or steadfast meditation, whatever you prefer – gives us more forms of control over our cognition. In particular, it gives us the ability to prevent ourselves from forming beliefs about p upon having considered evidence pertaining to p . Now suppose we find strong evidence that, say, the butler did it. We don't exercise our newly acquired ability and therefore automatically form the belief that the butler did it. Does the mere possession of this ability show that our mental state is not a genuine belief? I don't see why we should think that. By analogy, we shouldn't think that a JCA-TC2-style joint commitment is not a genuine group belief, just because the group could have refrained from forming it.

Here's a final worry about my response to the objection from voluntariness. My proposal holds that the joint commitments involved in group belief – i.e., those satisfying TC2 – are involuntary. But, one might object, this is confused. Beliefs are involuntary, but commitments are voluntary. If it isn't voluntary, it isn't a commitment.²⁹

First, I'm not convinced that the meaning of 'commitment' really is so clear and fixed as to preclude the possibility of commitments that are involuntary in the sense explained above. 'Commitment' is a philosophical term of art, the meaning of which must be partially fixed by stipulation. If there is good theoretical use for a category of involuntary commitments – which I suggest there is, if we're looking to make sense of group belief – then nothing stands in the way of using 'commitment' accordingly.

Second, even if I'm wrong about this semantic point, this objection *supports* rather than undermines my proposal when properly understood. Let's retrace the dialectic so far. The goal was to develop an account of group belief and the proposal was to do this in terms of a group's jointly committing to emulate belief in a proposition. The objection under consideration in this section is that this proposal is wrongheaded, since a group's joint commitment is under direct voluntary control and belief never is. To respond to this, I argued that joint acceptances which satisfy TC2 are *not* under the group's direct voluntary control. The objection to this that we are now considering is that my JCA-TC2 isn't really a joint *commitment* account anymore, because the commitments involved in it are not under direct voluntary control as commitments must be by definition. But in view of the original objection from voluntariness, this is precisely what we should have hoped for! In order to make the joint commitments involved in group beliefs sufficiently like beliefs, they have to be involuntary. This is what I've accomplished by adding the TC2. All that remains, then, is a verbal quibble about whether we ought to call the resulting account a joint *commitment* account. Because of the reasons mentioned in the previous paragraph, I'm inclined to say we should. But if someone insists that we really shouldn't, I'm happy to replace 'commitment' by a different word.

²⁹ Thanks to Rik Peels for raising this objection.

I conclude that the objection from voluntariness shows the basic JCA to be wrong. However, it fails as an objection to JCA-TC2.

7 Conclusion

The upshot is that JCA – or at least a version of it – can be defended against two powerful objections that have been leveled against it. This, however, requires a significant modification to the basic JCA, namely the addition of a condition that guarantees that the joint commitment in question bears the right sort of relation to the seeming truth of the proposition accepted. I proposed TC2 can do this job.

The addition, however, comes at a cost. The scope of JCA-TC2 is considerably narrower than that of the original JCA. Lackey's two counterexamples to JCA can easily be multiplied. This shows that plenty of joint commitments to emulate belief lack the sort of truth-connection required for genuine belief. Such joint commitments do not generate group *beliefs* proper, although we may call them that in ordinary loose talk, where we don't always distinguish carefully between genuine beliefs and states that may resemble it or its expressions, such as acceptances, hopes, lies, bullshit, etc. In other words, group belief along the lines of JCA exists, but its occurrence in real-life groups is rarer than its defenders have hitherto realized.

Acknowledgements The author would like to thank Matteo Colombo, Dominik Klein, Jan Sprenger, and the audience at the OZSW conference in Rotterdam in 2013 for helpful discussions about the topic of this paper. Thanks are also due to Jennifer Lackey, Kelvin McQueen, Rik Peels, Emanuel Rutten, Arthur Schipper, René van Woudenberg, Jan Willem Wieland, and K. Brad Wray for detailed suggestions about previous versions of this paper. Two anonymous referees also deserve thanks for their insightful and constructive comments. Part of the work for this paper was supported by a Vidi grant (276-20-024) from the Dutch Research Council (NWO) and a grant from the Templeton World Charity Foundation (TWCF). The opinions expressed in this publication are those of the author, and do not necessarily reflect the views of NWO or the TWCF.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adler, J. (2002). *Belief's Own Ethics*. Cambridge, MA: MIT Press
- Alston, W. P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257–299
- Cohen, L. J. (1989). Belief and Acceptance. *Mind*, 98(391), 367–389
- Conee, E. & Feldman, R. (1998). The Generality Problem for Reliabilism. *Philosophical Studies*, 89(1), 1–29

- Dunn, J. (2021). Reliable Group Belief. *Synthese*, 198(23), 5653–5677
- Feldman, R. (2001). Voluntary Belief and Epistemic Evaluation. In Steup, M. (Ed.), *Knowledge, Truth, and Duty* (pp. 77–92). New York: Oxford University Press
- Frankfurt, H. (2005). *On Bullshit*. Princeton, NJ: Princeton University Press
- Gilbert, M. (1987). Modelling Collective Belief. *Synthese*, 73(1), 185–204
- Gilbert, M. (1989). *On Social Facts*. London: Routledge
- Gilbert, M. (1994). Remarks on Group Belief. In Schmitt, F. F. (Ed.), *Socializing Epistemology* (pp. 235–256). Lanham, MD: Rowman & Littlefield
- Gilbert, M. (2004). Collective Epistemology. *Episteme*, 1(2): 95–107
- Gilbert, M. (2006). *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. New York: Oxford University Press
- Gilbert, M. (2014). *Joint Commitment: How We Make the Social World*. New York: Oxford University Press
- Gilbert, M. & Pilchman, D. (2014). Belief, Acceptance, and What Happens in Groups. In: J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 189–212). New York: Oxford University Press
- Ginet, C. (2001). Deciding to Believe. In M. Steup (Ed.), *Knowledge, Truth and Duty* (pp. 63–76). New York: Oxford University Press
- Goldman, A. (2014). Social Process Reliabilism: Solving Justification Problems in Collective Epistemology. In J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 11–41). New York: Oxford University Press
- Greco, D. (2014). A Puzzle About Epistemic Akrasia. *Philosophical Studies*, 167(2), 201–219
- Hakli, R. (2006). Group Beliefs and the Distinction Between Belief and Acceptance. *Cognitive Systems Research*, 7(2–3), 286–297
- Horowitz, S. (2014). Epistemic Akrasia. *Noûs*, 48(4), 718–744
- Lackey, J. (2020). Group Belief: Lessons from Lies and Bullshit. *Aristotelian Society Supplementary Volume*, 94(1), 185–208
- List, C., & Pettit, P. (2011). *Group Agency*. New York: Oxford University Press
- Mathiesen, K. (2006). The Epistemic Features of Group Belief. *Episteme*, 2(3), 161–175
- Mathiesen, K. (2011). Can Groups Be Epistemic Agents? In Sirtes, D., Schmid, H. B., & Weber, M. (Eds.), *Collective Epistemology* (pp. 23–44). Frankfurt: Ontos Verlag
- McHugh, C. (2012). Epistemic Deontology and Voluntariness. *Erkenntnis*, 77(1), 65–94
- Meijers, A. (2002). Collective Agents and Cognitive Attitudes. *Protosociology*, 16, 70–85
- Montmarquet, J. (1986). The Voluntariness of Belief. *Analysis*, 46, 49–53
- Nickel, P. (2010). Voluntary Belief on a Reasonable Basis. *Philosophy and Phenomenological Research*, 81(2), 312–334
- Peels, R. (2014). Against Doxastic Compatibilism. *Philosophy and Phenomenological Research*, 89(3), 679–702
- Pettit, P. (2011). Groups With Minds of Their Own. In Goldman, A. & Whitcomb, D. (Eds.), *Social Epistemology: Essential Readings* (pp. 242–268). New York: Oxford University Press
- Rolin, K. (2008). Science as Collective Knowledge. *Cognitive Systems Research*, 9(1–2), 115–124
- Schmitt, F. F. (1994). The Justification of Group Belief. In Schmitt, F. F. (Ed.), *Socializing Epistemology* (pp. 257–287). Lanham, MD: Rowman & Littlefield
- Schmitt, F.F. (2014). Group Belief and Acceptance. In Chant, S. R., Hindriks, F., & Preyer, G. (Eds.). *From Individual to Collective Intentionality: New Essays* (pp. 61–88). Oxford: Oxford University Press
- Shah, N. & Velleman, D. (2005). Doxastic Deliberation. *Philosophical Review*, 114, 497–534
- Steup, M. (2012). Belief Control and Intentionality. *Synthese*, 188(2), 145–163
- Sunstein, C. R. (2002). The Law of Group Polarization. *Journal of Political Philosophy*, 10(2), 175–195
- Tollefsen, D. (2002). Challenging Epistemic Individualism. *Protosociology*, 16, 86–117
- Tollefsen, D. (2003). Rejecting Rejectionism. *Protosociology*, 18: 389–405
- Tollefsen, D. (2015). *Groups as Agents*. Cambridge: Polity Press
- Tuomela, R. (1992). Group Beliefs. *Synthese*, 91(3), 285–318
- Tuomela, R. (2004). Group Knowledge Analyzed. *Episteme*, 1(2): 109–127
- Velleman, D. (2000). On the Aim of Belief. In *The Possibility of Practical Reason* (pp. 244–281). New York: Oxford University Press, 2000
- Wedgwood, R. (2002). The Aim of Belief. *Philosophical Perspectives*, 16, 267–297
- Williams, B. (1973). Deciding to Believe. In *Problems of the Self* (pp. 136–151). Cambridge: Cambridge University Press
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press

Wray, K. B. (2001). Collective Belief and Acceptance. *Synthese*, 129, 319–333

Wray, K.B. (2007). Who Has Scientific Knowledge? *Social Epistemology*, 21(3): 337–437

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.