

[Forthcoming in *Ratio*]

GROUP PEER DISAGREEMENT

J. Adam Carter

Abstract. A popular view in mainstream social epistemology maintains that, in the face of a revealed *peer disagreement* over p , neither party should remain just as confident *vis-a-vis* p as she initially was. This ‘conciliatory’ insight has been defended with regard to *individual* epistemic peers. However, to the extent that (non-summativist) *groups* are candidates for group knowledge and beliefs, we should expect groups (no less than individuals) to be in the market for disagreements. The aim here will be to carve out and explore an extension of the conciliatory insight from individual peer disagreement to *group peer disagreement*; in doing so, I’ll raise and address three key problems that face any plausible defence of such a constraint.

Suppose that you believe that p , and—as you just now find out—I believe not- p ; prior to our disagreement, we each reckoned the other to be equally likely to be right about whether p —*viz.*, we took each other to be *epistemic peers* on the matter¹. A widespread thought in recent social epistemology is that: at least in some circumstances like the one just described, a requirement of epistemic rationality is that both parties revise their respective p -beliefs in some fashion². Call this the *conciliatory insight*.

Conciliatory insight In a revealed peer disagreement over p , each individual should give at least some weight to her peer’s attitude such that neither is justified in staying exactly as confident as she initially was regarding whether p ³.

¹Epistemic peerhood is typically thought to involve both a cognitive and an evidential condition; as Lackey puts it, say someone is your epistemic peer, *vis-a-vis* p (or, more carefully: *vis-a-vis* the matter of whether p) just when (roughly) that individual and you are on a cognitive and evidential par, as your cognitive abilities and evidence bear on the matter of whether p . See here Jennifer Lackey, ‘A justificationist view of disagreement’s epistemic significance’, In *Proceedings of the XXII World Congress of Philosophy*, 53 (2008), pp. 145-154. Cf. Earl Conee, ‘Peerage’, *Episteme* 6, (2009), pp. 313-323.

²For some representative defences of this kind of position, see Adam Elga, ‘Reflection and disagreement’, *Nous*, 41(3) (2007), pp. 478-502; David Christensen, ‘Epistemology of Disagreement: The Good News’, *The Philosophical Review*, 116 (2007), pp. 187-217; and Richard Feldman, ‘Reasonable religious disagreements’, In Louise Antony (ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular*, Oxford University Press (2007), pp. 194-214. Cf. Tom Kelly, ‘The Epistemic Significance of Disagreement’, *Oxford Studies in Epistemology*, 1, (2005), pp. 167-196 and Peter van Inwagen, ‘It Is Wrong, Everywhere, Always, for Anyone, to Believe Anything upon Insufficient Evidence’, In Jeff Jordan & Daniel Howard-Snyder (eds.), *Faith, Freedom and Rationality*. Savage, Maryland: Rowman and Littlefield (1996), pp. 137-154, for notable opposition.

³I take this characterisation from Ballantyne & Coffman, ‘Conciliationism and Uniqueness’, *Australasian Journal of Philosophy*, 90(4) (2012), pp. 657-670.

The conciliatory insight has been motivated exclusively at the individual level—as a rational constraint on individual epistemic subjects. Can we make sense of the conciliatory insight at the group level? There’s no obvious barrier to doing so on a simple *summativist* conception of a group, according to which the attribution to a group of an epistemic state is just shorthand for an attribution to a majority of individuals an epistemic state. On the simple summativist proposal, we might firstly define epistemic peers summatively by saying that Group A and Group B are epistemic peer groups just in case enough individuals in group A are peers with enough individuals in group B. The conciliatory insight, applied to a summativist conception of group disagreement, reduces to the thought that the individuals in group A who recognise peer disagreement among individuals in B should revise their p-beliefs. Notice, though, that construed this way, the conciliatory insight remains ultimately a rational mandate on individuals.

But as Jennifer Lackey⁴ has remarked, ‘A fairly common view in current work in collective epistemology is that groups can have knowledge that not a single one of its members possesses’⁵. Consider, for instance—and following here Hutchins’⁶ classic example—of the crew of a ship navigating its way to port; different crew are responsible for charting different landmarks⁷. As Lackey puts it:

In such a case, the ship’s behavior as it safely travels into the port is clearly well-informed and deliberate, leading to the conclusion that there is collective knowledge present. More precisely, it is said that the crew as a whole knows, for instance, that

⁴ Jennifer Lackey, ‘Socially Extended Knowledge’, forthcoming in *Philosophical Issues*, special issue on *Extended Knowledge*, eds. Carter, J.A., Kallestrup, J., Palermos, O. and Pritchard, D. (2014), p. 1.

⁵Note that Lackey herself does not positively espouse Hutchins’ non-summativist insight with respect to group knowledge. For some notable proponents of the non-summativist insight, see here, along with Edwin Hutchins, *Cognition in the Wild*, (Cambridge, MIT Press, 1995); Ronald Giere, ‘Discussion Note: Distributed Cognition in Epistemic Cultures’, *Philosophy of Science*, 69, (2002), pp. 637-644; Karin Knorr-Cetina, *Epistemic Cultures: How the Sciences Make Knowledge*, (Harvard: Harvard University Press, 1999); Christine Halverson, *Inside the Cognitive Workplace: New Technology and Air Traffic Control* (Doctoral dissertation, University of California, San Diego, 1995); Margaret Gilbert, ‘Belief and Acceptance as Features of Groups’, *Protosociology*, 16 (2002), pp. 35-69; Margaret Gilbert, *Joint Commitment: How We Make the Social World*, (Oxford: Oxford University Press, 2013); Fred Stoutland, ‘The Ontology of Social Agency’, *Analyse & Kritik*, 30, (2008), pp. 533-551. Raimo Tuomela, *The Importance of Us: A Philosophical Study of Basic Social Notions*, (Stanford: Stanford University Press, 1995); and Deborah Tollefsen, ‘Collective Epistemic Agency’, *Southwest Philosophy Review*, 20, (2004), pp. 55-66. Thanks to an anonymous referee for suggesting clarification here.

⁶ Edwin Hutchins, *Cognition in the Wild*.

⁷Cf. Minsky’s (1988) remark that ‘each brain contains hundreds of different types of machines, interconnected in specific ways which predestine that brain to become a large, diverse society of partially specialized agencies’.

they are traveling north at 12 miles per hour, or that the ship itself knows this, even though no single crew member does⁸.

Call the position here *group knowledge*.

Group knowledge A group, G, can know that *p* even when not a single member of G knows that *p*.

The kind of groups that are candidates for (non-summative) group knowledge, in this sense, will not be *mere* collections of individuals (e.g. a crowd in New York's Rockefeller square); rather, the group's collective behaviour must be sufficiently integrated (e.g. expert panels, juries), and demarcated by a common aim⁹. Furthermore, such groups must generate *cognitive* outputs, in the form of representational states¹⁰.

It stands to reason that any system that generates cognitive outputs in the form of representational states can both represent correctly as well as misrepresent. Thus, the conditions that would make possible group knowledge make possible in equal measure group fallibility: where a group's cognitive output is *p*, even though not-*p* is the case. But given that the groups that can attain group knowledge are fallible, *group knowledge* implies what we can call *group disagreement*:

Group disagreement For two groups, G_1 and G_2 , it is possible that G_1 represents *p* and G_2 not-*p*.

Just as not all *individual* epistemic subjects are equally likely to get it right, the same goes for groups; in Hutchins' classic case of the ship crew, we can just as easily imagine a second ship crew of novices, less likely to navigate the ship to port. But sometimes, groups might be on a kind of epistemic par prior to a disagreement. Consider an idealised twist on Hutchins case:

DUELING SHIPS: Call the group constituting the ship crew on Hutchins' ship G_1 . Let G_2 be an equally experienced and able crew, all of whom have received training at the same naval academy as the crew in G_1 . (Assume moreover that this is common knowledge). Suppose G_1 and G_2 are navigating their respective ships to the port; let *p* be a proposition that if G_1 continues its course at 15 knots, with G_2 continuing its course at 20 knots, G_1 will collide with G_2 . G_1 , as a group, takes *p* to be false; G_2 , however, communicates to G_1 G_2 's position that *p* is true.

⁸ For a similar kind of case involving an airline cockpit, see Edwin Hutchins & Tove Klausen 'Distributed Cognition in an Airline Cockpit', *Cognition and Communication at Work*, (1996), pp. 15-34. See also Eric Winsberg, Bryce Huebner, & Rebecca Kukla, 'Accountability and values in radically collaborative research', *Studies in History and Philosophy of Science Part A*, 46, (2014), pp. 16-23, for a case involving collective scientific collaboration.

⁹ See here Gilbert, *Joint Commitment*. See also Christian List, 'Distributed Cognition: A Perspective from Social Choice Theory', in *Conferences on New Political Economy*, 25, (2008), pp. 285-308 and Christian List & Philip Pettit, 'Group Agency and Supervenience', in *Southern Journal of Philosophy*, 44, (2006), pp. 85-105, and Christian List & Philip Pettit 'Aggregating Sets of Judgements: An Impossibility Result', *Economics and Philosophy* 18, (2002), pp. 89-110

¹⁰ See here Christian List, 'Distributed Cognition: A Perspective from Social Choice Theory'.

Of course, given the significant *practical* stakes here, it's plausible that, purely for purposes of risk-aversion, both ships should (for instance) take immediate action to slow to a halt and wait for further information¹¹. But the 'should' here is one of *practical* rationality. Abstracting away from what the best practical policy is in such a case, it is also appropriate to ask whether, from a purely epistemic point of view, it is permissible for G_1 and G_2 to continue (after the discovered disagreement) to maintain, to the same extent as before, their 'cognitive outputs' of p and not- p , respectively.

It seems plain enough that the considerations that motivate the conciliatory insight, at the individual level, motivate a similar stance at the collective level, when conflicting cognitive outputs are reached (and discovered) by recognised peer groups¹². Thus, two views with notable support in social epistemology—the conciliatory insight, and the *group knowledge* thesis, seem to imply a kind of *group conciliatory insight*:

(Group) Conciliatory Insight In a revealed group-peer disagreement over p , each group should give at least some weight to the other (peer) group's position such that neither group is justified in staying exactly as confident as before regarding whether p ¹³.

As it turns out, though, the group conciliatory insight raises a nest of puzzles. For instance, in the DUELING SHIPS example, the group conciliatory insight mandates that each group revise its belief about whether p . But how exactly is this supposed to work? More specifically: what *must happen* for this requirement to be met?

Group Peerhood Disagreement: Three Problems

It should be emphasised that it would be puzzling at best if those who endorse the group knowledge thesis in collectivist epistemology are somehow preordained to avow an anti-conciliatory, 'steadfast' line in cases like DUELING SHIPS, where discovered peer-group disagreement is in place. But there is at least no problem in

¹¹Arguably, such a decision would be in alignment with what is called in health and environmental decision making, the *precautionary principle*—according to which, when threats to certain damages are great, certain preventative actions (e.g. in this case, slowing the ship) are warranted *even if* the threat to the damage is not one for which evidence is conclusive. For an overview, see here Martin Peterson, 'Should the Precautionary Principle Guide our Actions or our Beliefs?' *Journal of Medical Ethics* 33, (2007), pp. 5-10.

¹²As an anonymous referee has helpfully pointed out, this suggestion presupposes a kind of epistemic agency for (non-summativist) groups which would be precluded by Bryce Huebner's recent homoncular functionalist account of group belief, as outlined in Bryce Huebner, *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*, (Oxford: Oxford University Press, 2014).

¹³I take this characterization, with slight modification, from Ballantyne & Coffman, 'Conciliationism and Uniqueness'.

making sense of what steadfastness, for groups, would *involve-viz.*, the status quo. Group conciliatorism, by contrast, is much more perplexing on this score. Let's turn now to three problems that pose *prima facie* hurdles to vindicating the group conciliation insight.

The Doxastic Problem

A hallmark of individualist epistemology is the claim that propositional knowledge entails both truth *and belief*, (along with further conditions to the effect that known beliefs must be justified, non-accidentally correct, etc.). Presumably, then, true cognitive outputs issued by groups, in order to be knowledge, must be beliefs¹⁴. But the move from group knowledge to group belief is controversial. Hutchins' crew know that *p*; but does Hutchins' crew *believe* that *p*? As Hakli¹⁵ remarks:

One reason for this may be that the word 'knowledge' seems less tied to being a mental state [than belief]. In everyday language use, it is sometimes said that there is knowledge in books or libraries, for example, even though there is no agent that has the knowledge...Consequently, it might be easier to literally attribute knowledge to groups than to literally attribute beliefs to groups.

The debate between *rejectionists* (e.g. Meijers¹⁶, Wray¹⁷, Hakli¹⁸) and what have been dubbed *believers*: (Gilbert¹⁹, Tollefsen²⁰) centres around the issue of whether groups can have *bona fide* beliefs, or merely *acceptances*²¹. A problem for the group conciliatory insight materialises: if rejectionists are right, then regardless of whether the group knowledge thesis is correct, beliefs cannot be properly attributed to groups (in the non-summative sense)²². Obviously, groups can revise beliefs only provided

¹⁴ For a recent challenge to mainstream thinking on this point, see Blake Myers-Schulz & Eric Schwitzgebel, 'Knowing that P without Believing that P', *Nous* 47, (2013), pp. 371-384.

¹⁵ Raul Hakli, 'On the Possibility of Group Knowledge without Belief', *Social Epistemology*, 21, (2007), pp. 251.

¹⁶ Anthonie Meijers, 'Believing and Accepting as a Group', in *Belief, Cognition and the Will*, ed. A. Meijers (Tilburg; Tilburg University Press, 1999), pp. 59-71.

¹⁷ K. Brad Wray, 'Collective Belief and Acceptance', *Synthese*, 129, (2001), pp. 319-333.

¹⁸ Hakli, 'On the Possibility of Group Knowledge without Belief'.

¹⁹ Margaret Gilbert, 'Belief and Acceptance as Features of Groups', *Protosociology* 16, (2002), pp. 35-69.

²⁰ Tollefsen, 'Collective Epistemic Agency'.

²¹ For a helpful overview of some of the characteristic distinctions, see Pascal Engel, 'Believing, Holding True, and Accepting', *Philosophical Explorations* 1, (1998), pp. 140-151.

²² See Hakli, 'On the Possibility of Group Knowledge without Belief', for a rejectionist line on group knowledge to the effect that: group knowledge is justified, true acceptance. For an overview, see Kay Mathiesen, 'Introduction to Special Issue of Social Epistemology on Collective Knowledge and Collective Knowers', *Social Epistemology* 21, (2007), pp. 209-216.

that groups have beliefs to revise. Groups can't be rationally required to revise what they don't have²³. Therefore, it looks like if rejectionists are right, the group conciliatory insight cannot be vindicated. Call this the *doxastic problem*.

The Discharge Problem

Suppose, for the sake of argument, that the *doxastic problem* can be met and that groups can have beliefs no less than knowledge. A general normative problem arises, one that will crop up for any particular account of group belief²⁴; the problem, in short, is how to make sense of a group obligation to undertake doxastic revision? More specifically, how can we make sense of how this rational obligation might be discharged from a group G to members? It's not obvious after all what the members of G are meant to do in order that G does not count as (and in violation of the group conciliatory insight) staying exactly as confident as before in its original cognitive output.

To bring this problem into sharp relief, it is helpful to consider a plausible constraint on any group action, which Lackey calls the, Group/Member Action Principle:

Group/Member Action Principle (GMAP) For every group, G, and act, α , G performs α only if at least one member of G performs some act or other that causally contributes to α .²⁵

What GMAP preserves, Lackey says, 'is that while group action *cannot occur independently of its members*, it can go beyond what any of them do individually²⁶' (*Ibid.*, 7).

What GMAP implies, thus, is that, in our DUELING SHIPS case, G₁ can revise its belief that p only if at least one member of G₁ performs some act that causally contributes to the group's collectively revising its belief. *But what would such an act be?* Of course, such a requirement borders on incoherence were we to suppose that G₁ believes p , and G₂ not- p , but that no *individual* in G₁ is aware that G₂, as a group, believes not- p ²⁷. What this reveals is that a precondition for the kind of *recognised*

²³For a discussion of the platitude that ought implies can, extended to groups and group abilities, see Holly Lawford-Smith, 'The Feasibility of Collectives' Actions', *Australasian Journal of Philosophy*, 90, (2012), p. 453.

²⁴In §3 I take up a problem that arises more directly for the nature of group belief in terms of undertaking certain commitments.

²⁵ Lackey, 'Socially Extended Knowledge', p. 7. For a similar statement, see Pekka Makela and Raimo Tuomela 'Group Action and Group Responsibility', *Protosociology*, 16 (2002), p. 198.

²⁶My italics.

²⁷To be clear, I am not suggesting that it is incoherent that group members not be aware when group disagreements are in place. It is after all not incoherent that an individual be a member of a group that believes p , even though the individual is not aware that the group believes p (e.g. see Huebner (2014)). Rather, the point is the following: If GMAP is correct, then G₁ can revise its belief that p only if at least one member of G₁ performs some act that causally contributes to the group's collectively revising its belief. Now, what I suggest is borderlines on incoherence is that this rational requirement be in

group peer disagreement that—by reference to the conciliatory insight—would motivate doxastic revision, is the following: that the relevant peer disagreement be recognised by at least one individual in the group²⁸.

Suppose this condition is satisfied. Even so, with reference to GMAP, a kind of discharge problem raises its head: what can any individual be *obligated* to do to causally contribute to the group’s revising its belief? Suppose Stan is a member of Hutchins’ crew, and Stan is aware both that his crew— G_1 collectively believe p , and further, that G_2 (a recognised peer group) believes not- p . What, with reference to the group conciliatory insight, is Stan’s obligation? We might be tempted to say that his obligation is to act so as to revise G_1 ’s collective belief that p . The background rationale might be that: ought distributes over conjunction in such a way that: if G_1 has an obligation to ϕ , then so does each of its members.

A moment’s reflection shows that Stan’s obligation can’t be accounted for in this way. Just suppose, to borrow a simple example from Lawford-Smith²⁹ that A, B, C and D are four equal-partner employees of a moving company—ABCD Moving—and that, while they are working a job, they accidentally drop a piano so that a child is trapped. ABCD Moving has an obligation to remove the piano. But if the ought binding the group distributes over the conjunction of its members such that A thereby has an obligation to remove the piano, then a simple problem quickly emerges³⁰: just suppose that A cannot lift the piano alone, and B, C, and D do not help. A cannot then be obligated to move the piano³¹.

What this means, though, is that if Stan’s obligation—as circumscribed by the group conciliatory insight—is to revise G_1 ’s belief that p , it *can’t be simply* because (as the group conciliatory insight insists) G_1 has this obligation.

This point, to be clear, does not entail either that the group conciliatory insight is false nor that ought doesn’t distribute over conjunction. (After all, ought might

place for G_1 even in a situation where no individual in G_1 is at least aware of the group belief disagreement. This is, in effect, a point in line with ‘ought implies can’. The same suggestion holds at the individual level, in that, the epistemological question of what is rationally required in the face of peer disagreement is always a question about *revealed* or *recognised* peer disagreements. Thanks to an anonymous referee for requesting clarification on this point.

²⁸Thus, there is no meaningful sense in which the recognition requirement undergirding the rational requirement of belief revision will be satisfied in a situation where the individuals of a group are not aware what the group’s cognitive output is. See here, for instance, Huebner, *Macro cognition*.

²⁹ Lawford-Smith, ‘The Feasibility of Collectives’ Actions’.

³⁰ Cf. Frank Jackson & Robert Pargetter, ‘Oughts, Options, and Actualism’, *The Philosophical Review* 95, (1986), pp. 233-255.

³¹ See Lawford-Smith, ‘The Feasibility of Collectives’ Actions’, p. 456. For an earlier presentation of this case, see Raimo Tuomela & Kaarlo Miller ‘We-intentions’, *Philosophical Studies* 53 (1988), pp. 367-389.

distribute over conjunction in ways other than the simple way just considered). Rather, and put carefully, the point is that whatever it is that Stan is (by reference to the group conciliatory insight) rationally obligated to do, it isn't something that can be explained so simply as to say that Stan is obligated to revise G_1 's collective belief that p . Some further sophistication is thus needed to account for how rational obligation for G_1 to revise its belief that p might be discharged to its members. Call this the *discharge problem*.

The Problem of Normative Conflict

Let's suppose that the discharge problem is addressed, such that Stan has some obligation O , and that he has O in virtue of G_1 's obligation to revise its belief that p . By GMAP, Stan's performing O must causally contribute to G_1 revising its belief. We might initially suppose that O is the obligation to inform the other crew members that G_2 believes not- p . But this obligation is too weakly construed. To see why, simply imagine a modified case where each member of G_1 is like Stan in so far as each member in G_1 learns that G_2 believes not- p . If Stan's obligation, *qua* member of G_1 , is *merely* to inform the others that G_2 believes not- p , then consider that the other members of G_1 , $M_1...M_n$, will also have discharged their respective obligations simply by *telling* other G_1 members that G_2 believes not- p . But it obviously doesn't follow from all members $M_1...M_n$ in G_1 being *aware* that G_2 believes not- p that G_1 has—through $M_1...M_n$'s awareness—satisfied its conciliatory obligation to no longer maintain its belief to the same extent. It would remain an open question whether, after such awareness, G_1 then actually counts as revising its belief that p ³².

Thus, whatever obligation individuals in G_1 have in virtue of G_1 as a group being rationally required to revise its p -belief, the obligation is stronger than merely reporting to others in the group *that* G_2 's belief is in conflict with G_1 's belief. In order to (as per GMAP) *causally contribute* to G_1 's revising its belief, then, Stan must perform some action that is a part of a causal chain that results not *just* in the members of G_1 appreciating *that* G_1 's group belief conflicts with G_2 's group belief—but rather, Stan must perform some action that is part of a causal chain that results in G_1 no longer holding p to the same extent as before.

Plausibly, such an action will at least involve the following: encouraging the other members of the crew to act in ways that will jointly cause G_1 to revise, groupally, its previous doxastic position *vis-a-vis* p , even if just by degree³³.

³²Compare: in the individual case, one issue is whether peer disagreement is revealed. Another is whether each party to the revealed peer disagreement revises her belief in the face of the revelation. The revelation is not itself a kind of revision.

³³See Elga, 'Reflection and disagreement', for a proposal according to which the rational stance to take is to split the difference. Feldman, 'Reasonable religious disagreements', by contrast recommends agnosticism with respect to p .

But here a problem of normative conflict emerges. The obligation just outlined—to act in ways that will cause members of G_1 to jointly act so as to revise G_1 's previous doxastic position *vis-a-vis* p , seems to directly conflict with an independent obligation that each member of G_1 has in virtue of being a member of a group that collectively believes p .

To appreciate this point, consider here Gilbert's classic 'plural subject' account of group belief³⁴, according to which a group belief is formed when the members jointly accept a view³⁵. Joint acceptance is then understood as a kind of *conditional commitment*: when group members conditionally commit to accept that p , each is committed to accepting that p provided the others do³⁶.

So, what is one committing to when committing to accept that p provided others do? Gilbert says:

One who participates in joint acceptance of p thereby accepts an obligation to do what he can to bring it about that any joint endeavors... among the members of G be conducted on the assumption that p is true. He is entitled to expect others' support in bringing this about...if one does believe something that is inconsistent with p , one is required at least not to express that belief baldly³⁷. (Gilbert 1989, pp. 306-7)

In short, a group believes that p only if its members act as if p ; and in being a part of a group that believes that p , one is obligated to act as if p , and the obligation to act as if p includes an obligation to suppress one's doubts *vis-a-vis* p ³⁸.

This suppression requirement of course looks to be in stark tension with the thought that the group conciliatory insight could ever be satisfied³⁹; after all, how can (for

³⁴ E.g. Gilbert, 'Belief and Acceptance as Features of Groups'. Note that Gilbert has moved toward an exclusively 'joint commitment approach', and no longer relies on the language of 'plural subject'. Thanks to an anonymous referee for raising this point.

³⁵ Note that the problem of normative conflict sketched in this section is specific to joint commitment accounts of group belief. Accounts of group belief that do not rely on a joint commitment condition (as Gilbert does) do not face this particular problem. Thanks to an anonymous referee for suggesting clarification here.

³⁶ See Gilbert (1996, 204-05; cf. Gilbert (2013)); see also here for a helpful overview Mathiesen (2006, 168). Note that joint commitments are fundamentally group, not individual, commitments.

³⁷ My italics. Much of this passage is quoted also in Lackey (2012).

³⁸ Cf. Mathiesen 2006 for resistance on this point, by appealing to group epistemic obligations. I am sympathetic to Mathiesen's argument and will be challenging the simple suppression interpretation of Gilbert's requirement in §3.

³⁹ At this point, the following point could be raised: why not simply abandon the suppression requirement? After all, as this line of thought goes, the tension in question evaporates if we can simply reject out of hand Gilbert's suppression requirement. This move, however, cannot be pursued so easily. Something very much like the suppression requirement is going to naturally fall out of any

instance) Stan (*ala* GMAP) causally contribute to bringing the members of G_1 to revise their belief that p if Stan is obligated, qua part of a group that believes that p , to suppress his doubts about p ?

One might try to find a way out of this jam by exploiting a loophole Gilbert (2004, 104) offers—whereby one is allowed to express one’s doubts about the veracity of the group belief in a ‘personal’ capacity, provided this expression is not *as* a member of the group⁴⁰. But, it’s hard to see exactly what kinds of expressions of doubts (with reference to this loophole) will remain permissible. Mathiesen (2006, 169) nicely captures the worry:

Does it mean that the member must make clear that she recognizes that what she is about to say is not the official group view, but nevertheless she intends to question that view and intends that the others engage with her on this topic as group members? Or does it mean that the member must make clear that she is merely expressing a ‘personal belief’, which neither she nor the other members of the group should see as part of the group belief forming and revising process?

One thing should be clear: any attempt for Stan to discharge his obligation (via the group conciliatory insight) by acting to causally contribute to G_1 ’s revising its collective belief that p , can’t (with reference to Stan’s commitment and corresponding obligation to accepting that p) be *via* the expression of doubts as a member of the group, but only as personal doubts (whatever this would mean).

But now an even deeper conflict appears: it seems that it is specifically in virtue of Stan’s membership in a group that believes p that (by the group conciliatory insight) he is obligated to act to cause the group to jointly revise its belief that p . But this obligation is an obligation Stan has because he is member of the group—and so his expressions of doubt and actions that would causally contribute to G_1 ’s joint revision of its belief that p cannot be merely personal. They are, by the group conciliatory insight, precisely his obligations as a member of the group.

construal of non-summativist group belief as conditional commitment; without individuals in a group committed to acting as if p , it’s not clear how the group can be said to believe that p . I do not think the problem turns out to be intractable, as will be suggested in the positive proposal I sketch in the §3.2. Nonetheless, that something like suppression seems to fall out of the kind of conditional acceptance to act as if p is true that is an element of group belief, it suffices for the present purposes to see that avoiding the tension is not so easy as simply dropping the suppression requirement. Thanks to Jennifer Lackey for helpful suggestions for clarification on this point.

⁴⁰I should admit, however, that I do not here have a principled explanation of what features of acting as if p are relevant to not undermining joint commitment. In particular, the philosophical problem is: what features of acting as if p are relevant to undermining joint commitment, and *why*? In Hutchins’ case, for instance, it seems clear enough that blowing a signalling horn that all crew recognise would be blown only if p is false, is an instance of not acting as if p , in a way that would be relevant to undermining joint commitment. But cases of, for instance, flipping a switch, become more complicated. I do not at present have worked out principled conditions here, though I think any such principled account would need to be informed by context-relevant details of cases. Thanks to an anonymous referee for raising this issue.

Group Peer Disagreement and the Conciliatory Insight Revisited

An Awkward Result

The *doxastic problem*, the *discharge problem* and the *problem of normative conflict*, taken together, spell trouble for the prospects of extending the conciliatory insight from individual peer disagreements to group peer disagreements.

Does this mean that the group conciliatory insight should be rejected? It does seem that way. But what is the alternative? As Tom Kelly, points out, there are actually a spectrum of positions that can be endorsed by one who rejects the conciliatory position⁴¹. One alternative incompatible with the group conciliatory insight is what we can call *Moderate Group Steadfastness*, a position that denies the conciliatory claim that the fact of revealed peer disagreement always rationally requires some level of belief revision.

Moderate (Group) Steadfastness In at least some cases of revealed group peer disagreement over p , it can be reasonable to stick to one's guns; in other cases, some non-negligible belief revision is required.

The 'non-negligible' belief revision that is rationally required by moderate steadfastness views can be considerably less demanding than, for instance, Elga's equal weight position. But moderate steadfast views will insist that in some cases of revealed peer disagreement, at least *some* weight must be given to the peer's position.

Awkwardly, it looks like the three problems outlined in §2 will raise problems not only for the group conciliatory insight, but also, for *moderate group steadfastness*, in so far as the moderate position allows that there are at least some instances in which group peer disagreement will rationally require some (non-negligible) group doxastic revision. But this means that if we reject the group conciliatory insight on the basis of the reasons outlined in §2, then what we are left with is not a moderate group steadfastness, but rather, with a stronger position which we can call *strong group steadfastness*.

Strong (Group) Steadfastness In all cases of revealed group peer disagreement over p , it is reasonable to stick to one's guns—such that it is not unreasonable for both parties to remain as confident as they initially were regarding whether p

Not even the most ardent critics of the conciliatory view, at the individual level, want to take on board the strong steadfast view that insists it is always rationally permissible for both parties to hold their guns in all circumstances of revealed group peer disagreement.

⁴¹See here Thomas Kelly, 'Peer Disagreement and Higher Order Evidence', in Richard Feldman and Ted Warfield (eds.), *Disagreement* (Oxford: OUP, 2010), §1.

Strong Group Steadfastness is thus an unwelcome consequence. But we either attempt to find *some way* to make sense of the requirement that groups rationally revise their beliefs in the face of group peer disagreement (despite the worries raised in §2) or accept the implausibly strong group steadfastness position. The former approach looks more promising.

Action, Risk and Representation

While it's evident that the *doxastic problem*, the *discharge problem* and the *problem of normative conflict* will significantly constrict how it is that we might compellingly spell out how the group conciliatory insight could be respected, in practice, it's no less evident that the strong group steadfastness alternative simply won't work in theory. In this section, I want to sketch what I think is a potential way out of this impasse—a limited defence of a way to hold on to group conciliationism in light of the problems advanced.

To this end, consider first the doxastic problem. The doxastic problem in short says that groups don't have the sort of thing that can be revised—beliefs; rejectionists offer a number of reasons for thinking that groups have acceptances rather than beliefs⁴²; thus, it looks like the group conciliationism can be vindicated only if one is prepared to enter the debate between rejectionists and believers *vis-a-vis* group belief and demonstrate that rejectionists are wrong and believers are right.

But I think there is a sensible way to defend group conciliationism without this dialectical burden. The doxastic problem is a problem, after all, for group conciliationism only if we accept *ex ante* a certain picture about what doxastic revision would involve. This picture says that beliefs can be revised (perhaps, by lowering credences) and acceptances cannot.

However, there's room to resist things at this point. Consider here Preyer's⁴³ statement of an idea that undergirds the rejectionist thinking about the difference between belief and acceptance, as this bears on the matter of revision:

Beliefs are gradual (quantitative) and have subjective probabilities, acceptance is a matter of quality...They have a feature of all or nothing.

With this idea in play, Meijers (2002)—a leading proponent of the rejectionist position—reasons in the following way. Gilbert-style group beliefs involve conditional agreements, where such agreements are essentially an on-off matter. As Meijers puts it, 'We agree, or we do not', (*Ibid.*, 81). Thus, we should reject that groups have

⁴²See here Meijers 'Believing and Accepting as a Group', Gerhard Preyer, 'What is Wrong with Rejectionists?' In *Interpretation, Sprache und das Soziale: Philosophische Artikel*, (Frankfurt am Main: Humanities Online, 2003) and Kay Mathiesen, 'The Epistemic Features of Group Belief', *Episteme*, 2, (2006), pp. 161-175.

⁴³ Preyer, 'What is Wrong with Rejectionists?'

beliefs, because beliefs don't have such an all-or-nothing feature; rather, they are gradient.

There is some space though for the proponent of group conciliatorism to intercede here. Whilst agreement is indeed an on-or-off affair, there is a matter of flexibility with respect to the object of agreement, when that object is a representational attitude. We might agree as a group to endorse a p without qualification, or with certain qualifications, regarding just *how* we shall act as if p .

The ways we are obligated to act, as a result of this agreement, will be in part a function of any such qualification made. At one limit, where there is no qualification, we might agree to act (as Gilbert's view suggests) as if p (as a member of the group) in our talk and actions, no matter what doubts we harbour, no matter the circumstance.

Though we might agree to jointly accept p , with certain qualifications, and again, how we're obligated to act will be in part a function of such qualifications. We might for instance agree to accept p as a group and, in doing so, agree to rely on p in ways that track the practical stakes.

At this point, it will be useful to make a comparison to the case of an individual believer and a propositional attitude, α . There is a range of propositional attitudes $\alpha_1 \dots \alpha_n$ that I might have, *vis-a-vis* some proposition p , which correspond to my willingness to act as if p , given certain stakes⁴⁴.

At one end of the spectrum of propositional attitudes, for instance, we might envision the attitude of 'being certain that', where it is reasonable for me to act on my attitude of *being certain that* p no matter what the practical risk of being wrong. But we can also imagine attitudes toward the other end of the spectrum, which are reasonable to rely on only in circumstances where there is very little risk to being wrong, but *not* when such risks are great. For example, the attitude of *suspecting that* or *conjecturing that* p fits this profile.

Now, regardless of whether I am certain that p or merely suspect that p , in *both* cases I accept that p in so far as I am willing to rely on p across certain counterfactual circumstances. What differs dramatically is the circumstances.

Returning now to the group case; rather than to delimit the group's range of propositional attitudes *vis-a-vis* p to believing that p or accepting that p , we can (as in the individual case) envision a range of propositional attitudes $\alpha_1 \dots \alpha_n$ a group might have, *vis-a-vis* some proposition p , which correspond to the group's willingness to rely on p , given certain stakes.

⁴⁴For a much more detailed expression of this kind of idea, see J. Adam Carter, Benjamin Jarvis & Katherine Rubin, 'Belief without Credence', (2014), unpublished.

The agreement a group makes, in agreeing to accept p as a group, might thus be an agreement to rely on p *more or less*, depending on the circumstances. If for instance the group's epistemic position *vis-a-vis* p is weak, the group might be reasonable to agree as a group to accept p , in the following qualified way: the group agrees to rely on p only when there is little risk to being wrong. In such a situation, (and analogously with the case of an individual agent) we might say the group merely suspects or conjectures that p . Alternatively, a group might jointly agree (in light of a strong epistemic position) to accept p in a very different way: to rely on p even if there is great risk to being wrong. In such a situation, we might say the group is certain that p .

In this respect, we can see that—without engaging directly with the rejectionists about group belief—there is scope to argue that a group's taking up a representation attitude *by agreement* isn't (simply *qua* a product of agreement, *per se*) incompatible with a group adopting a *gradient* hierarchy of attitudes.

With this picture in mind, we have the resources to address the doxastic problem; a possibility for modelling group belief revision emerges, and which is compatible with the Gilbert-style thought that group belief is a matter of *agreement*. Just consider that a group can agree to accept p by agreeing to act as if p in the following qualified way: the group can agree to act as if p *across a range of circumstances that positively correlates with the group's epistemic position*, which is a position the individuals of the group agree to (and are thus obligated to) update.⁴⁵

A group that has agreed to accept p , and thus agreed to act as if p in the fashion just described, might accordingly (in the face of revealed group peer disagreement) downgrade its position from being almost sure that p to merely suspecting p in the following way: the group will have agreed that when evidence comes in against p , the group will move from acting as if p across most *all* practical environments, to acting as if p only in a narrow set of circumstances, where there is little practical risk of being wrong.

It should be clear how the model suggested here not only offers a way around the doxastic problem, but also, that it offers a simple solution to the problem of normative conflict. The problem of normative conflict emerged when it seemed that whatever one might do to cause the group to revise its belief would be strictly incompatible with a group member's obligations *qua* a member of a group that has agreed to act as if p is true. However, on the proposal offered here, speaking out when new evidence comes against the proposition accepted by the group is precisely what one agrees to do when agreeing to act as if p is true across a range of

⁴⁵Note that the norms governing acceptance here are fundamentally epistemic, in that what fixes the circumstances under which one which acts as if p , that correspond with different doxastic attitudes, is strength of epistemic position. On the difference between epistemic versus practical norms of acceptance, see Joëlle Proust, 'The Norms of Acceptance', *Philosophical Issues*, 22, (2012), pp. 316-333.

circumstances that positively correlates with the group's current (and updated) epistemic position, *vis-a-vis* p ⁴⁶. Thus, there is no conflict.

This leaves of course the discharge problem. I will not attempt here to answer the more general philosophical problem of what the conditions are for any member to incur an obligation as a part of a group. It should suffice at this point to point out that the proposal I've offered renders it much less mysterious what (by GMAP) one could do to cause the group to revise its collective doxastic stance towards p . The individual is obligated (as she acquires new evidence) to herself act as if p is true in a range of circumstances that aligns with the updated evidence, and to encourage others to act this way. For instance, when Stan learns that G_2 believes that not- p , Stan's obligation is to (for his part) limit the range of circumstances in which he will act as if p is true to those scenarios where the cost of being wrong is lower. He is obligated to encourage others to do the same. Since in the case of DUELING SHIPS, the practical cost is high, Stan is obligated to refrain from proceeding as if p is true, and to encourage others accordingly; in doing so, he will have (in alignment with GMAP) causally contributed to G_1 's downgrading its belief that p to a mere conjecture that p , one the group acts on only across a more narrow range of circumstances than before. Plausibly, this is precisely what Hutchins' ship ought to do in the face of group peer disagreement. Despite the initial hurdles, then, it looks like collectivist epistemology does in the end have the resources to embrace a kind of group conciliatory insight that parallels the widely held conciliatory insight in individualist epistemology. Given the appeal of this insight in individualist epistemology, this is good news⁴⁷.

Eidyn Research Centre, University of Edinburgh

5.04 Dugald Stewart Building, 3 Charles St., Edinburgh, EH8 9AD

j.adam.carter@ed.ac.uk

⁴⁶Note that the problem of normative conflict arise because the commitment each member of the group has, qua member of the group that believes p , is to act as if p . On such a construal, it seemed there was no room for a member of such a group to do any action that causally contributes to group belief revision, as such an action would appear to violate one's commitment to act as if p . However, the proposal I've sketched here has attempted to circumvent this tension by going beyond what Gilbert says in specifying how one counts as acting as if p (as would be established by the commitment).

⁴⁷I would like to thank an audience at Institute Jean Nicod, where this paper was presented in March 2014, for helpful feedback. Special thanks also to Emma C. Gordon, Jennifer Lackey, S. Orestis Palermos and Duncan Pritchard for comments on an earlier version of the paper. Thanks also to an anonymous referee at *Ratio* who offered a number of very helpful suggestions. This article was written as part of the AHRC-funded 'Extended Knowledge' (#AH/J011908/1) research project that is hosted by the University of Edinburgh's *Eidyn* research centre.