

GROUP KNOWLEDGE AND EPISTEMIC DEFEAT

J. ADAM CARTER

University of Edinburgh

If individual knowledge and justification can be vanquished by epistemic defeaters, then the same should go for group knowledge. Lackey (2014) has recently argued that one especially strong conception of group knowledge defended by Bird (2010) is incapable of explaining how it is that (group) knowledge is ever subject to ordinary mechanisms of epistemic defeat. Lackey takes it that her objections do not also apply to a more moderate articulation of group knowledge—one that is embraced widely in collective epistemology—and which she does not challenge. This paper argues that given certain background premises that are embraced by orthodox thinking in collective epistemology, the more moderate account of group knowledge cannot make sense of either psychological or normative epistemic defeaters. I conclude by offering some suggestions for how the more moderate proposal might avoid this result.

1. Introduction

Few philosophers are doing more these days to bridge the gap between mainstream (individualist) epistemology and collective epistemology than Jennifer Lackey,¹ and one of the most interesting applications of mainstream epistemology she has made to the epistemology of groups concerns the topic of epistemic defeat.

Epistemic defeaters for an individual's (would-be) knowledge and epistemic justification are reasonably well understood.² The very simple picture is this: *S*'s knowledge or justification that *p* can be defeated either *psychologically* or *norma-*

Contact: J. Adam Carter <j.adam.carter@ed.ac.uk>

1. See, along with other recent and forthcoming papers, especially Lackey (2012; 2014). Lackey's work on collective epistemology draws from recent work in individualist epistemology, in particular social epistemology. Her planned monograph *The Epistemology of Groups* is forthcoming with OUP.

2. For some sample comprehensive discussions of epistemic defeat, see especially Bergmann (1997), Lackey (1999; 2005), Pollock (1986). See also Sudduth (2008) for a contemporary overview.

tively, in virtue of S in the former case *possessing* some belief, q or, in the latter, being such that S (epistemically) *should* believe q , where q is a proposition the truth of which would either (i) count against the truth of S 's belief that p —this is called a *rebutting defeater*—or (ii) indicate that S 's belief that p was unreliably formed or sustained—this is called an *undercutting defeater*. So there can be rebutting or undercutting psychological and normative defeaters. Defeaters themselves, can be defeated by other defeaters. However, when they are not, (undefeated) defeaters vanquish knowledge and justification.

Here's a compelling and simple thought: if individual knowledge and justification can be vanquished by defeaters, then so too can group knowledge. In fact, it would be very odd indeed if group knowledge and justification were somehow *undefeatable*. (Two minds are not *that* much better than one!) The upshot of this point is that it is an *ex ante* constraint on any plausible account of group knowledge that it be in principle defeated through ordinary mechanisms of defeat; and if staying true to this constraint should come at the cost of implausible consequences, then all the worse for that theory of group knowledge.

With this kind of constraint in hand, Lackey (2014) has sought to show that Alexander Bird's somewhat radical account of group knowledge is not viable.³ I think Lackey is right, but that her argument does not go far enough. What she overlooks is that there is a way to argue reasonably straightforwardly to the conclusion that the more orthodox account of group knowledge (to which she doesn't level this criticism) has no viable way to make sense of epistemic defeat either, and (unless the problem can be resolved) this stands to be a very serious problem for orthodoxy in collective epistemology.

Here is the plan. In Section 2 I distinguish between the radical account of group knowledge that Lackey targets, and the more moderate view she doesn't. In Section 3, I outline just how it is Lackey thinks the radical view is hopeless in making sense of epistemic defeat. In Section 4 I argue that, at least on one very standard way of thinking about group belief (i.e., the joint commitment model), the more moderate view is just as hopeless in making sense of epistemic defeat as the more radical view. Section 5 then sketches some refinements to the joint commitment approach to group belief and shows how these refinements reveal some potential ways to save the more moderate account of group knowledge from the kinds of defeat-based objections raised in Section 3.

3. Lackey's argument from epistemic defeat is one of several arguments she offers against Bird's account of (radically) socially extended knowing. Another key element in Lackey's critique of Bird involves the relationship between knowledge and action. For the present purposes, I will be focusing entirely on her argument from defeat. Though it is important to keep in mind that Lackey's argument against Bird does not rest entirely on the argument from defeat.

2. Two Kinds of Socially Extended Knowledge

Let's now get clearer about the kind of group knowledge that will be at issue in what follows. Mainstream epistemology has traditionally concerned itself with the epistemic properties of the cognitive outputs of *individuals* (e.g. individual beliefs⁴). But we also attribute epistemic states such as belief and knowledge to groups; we ask what the FBI *knew* before 9/11 or what CERN knows about the Higgs-Boson and whether the jury was justified in its conclusion.⁵

Group knowledge is uncontroversial, *per se*, by the lights of mainstream epistemology, provided it is understood *summatively*. On a summativist conception of group knowledge, a group *A* knows that *p* iff most individuals in *A* know that *p*. Group knowledge, in this deflated sense, is just an aggregate of *individual* knowledge. A much more philosophically interesting conception of group knowledge is *non-summativist*. For the present purposes, it will be helpful to distinguish two very different non-summativist views, as Lackey (2014) herself does, along the lines of what each thesis epistemically *tolerates* of the members of group that collectively knows. First consider the version of the thesis that is widely endorsed in collective epistemology—call this *Socially Extended Knowledge-Moderate* (SEK-M):⁶

Socially Extended Knowledge-Moderate (SEK-M): A group, *G*, can know that *p* even when not a single individual member of *G* knows that *p*.⁷

A classic example, one which Lackey notes, that used to motivate (non-summativist) group knowledge is given by Hutchins (1995) in his famous case of a ship crew skillfully navigating an Iwo Jima class amphibious assault ship to port.

As Lackey summarises:

In such a case, the ship's behaviour 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 they are traveling north at 80 miles per hour, or that the ship itself knows this, even though no single crew member does. (Lackey 2014: 282)

4. For a clear sense of the project of analysing knowledge, understood along individual lines, see Ichikawa and Steup (2014).

5. For a comprehensive discussion of group knowledge attributions, see Lackey (2012).

6. For some approaches to group knowledge friendly to (SEK-M), see for instance Gilbert (1987; 2002; 2013), Knorr Cetina (2009), Mathiesen (2006), Tollefsen (2006), Tuomela (1995), Wray (2007).

7. One kind of case that proponents of distributed cognition are happy to embrace is Hutchins's 1995 case of ship crew skillfully navigating an Iwo Jima class amphibious assault ship to port. On Hutchins' view, the crew can know what no individual member knows.

Whereas (SEK-M) tolerates an absence of individual knowledge, it nonetheless allows for epistemic requirements weaker than knowledge of its members, viz., that individual members jointly accept that p or jointly commit to the proposition that p .⁸ There is, however, a more radical statement of group knowledge that Bird (2010) calls *social knowing*, which relaxes such a requirement *entirely*.

Socially Extended Knowledge-Radical (SEK-R): A group G can know that p even when not a single individual member of G is aware that p .⁹

As Alexander Bird (2010) sees it, all a social structure must do to be a candidate social knower is to have (in short) characteristic propositional outputs, and characteristic mechanisms that ensure the reliability of the propositional outputs, which are themselves the *inputs* for social actions or for social-cognitive structures, including the same structure that produces the output.¹⁰ Nowhere in this

8. Lackey (2014) herself uses a case of distributed cognition to make the point that (SEK-M) is a relatively standard view in collective epistemology. I follow her in using this case as demonstrative of (SEK-M). However, distributed-cognition-style cases aren't needed to make the point. Consider, for example, the following case, which achieves the same result—viz., that a group can know something non of the individual members knows—with the strictures of a joint-commitment model.

Committee: An art expert authentication committee has been called on to determine the veracity of a seller's claim that a particular pot dates to the late Jōmon period (circa 300 AD). The committee is comprised of three members, A, B and C, each highly skilled in thermoluminescence, a dating method for archaeological pottery that is used to distinguish genuine from fake antiquities. The committee members adhere to strict professional and scientific norms and (after pooling all evidence) reach a clear group position: that the pot is much older than expected. It dates to the early Jōmon period (approximately 14,000 BC). All three committee members, however, are young earth Creationist and so, in a private capacity, each believes the earth, and thus the pot, is less than 6,000 years old.

Provided that group knowledge involves group belief or acceptance (understood along joint-commitment lines) COMMITTEE is a case where the group, whose members adhere to the strictest scientific standards, knows something—that the pot dates to the early Jōmon period—even though each of its members fails to know because, due to their faith, each fails to believe the proposition in question.

9. While Lackey uses 'aware' to capture the laxity of the requirement, it's important to note that she is not using awareness in the sense that: if one is aware that p , then one knows that p —or else the radical view collapses into the moderate view. Rather, I think we could capture the spirit of Lackey's attempt to characterise the position Bird is committing himself to by reading 'not aware' that p as 'completely oblivious to' p . As we'll shortly see, Bird's view allows group knowledge in cases where all individuals who had any clue about the item of knowledge in question had *died*. In this sense, these individuals were not *aware* of (and were completely oblivious to) p .

10. Here I am following Lackey's (2014, p. 284-5) discussion of the key thrust of Bird's position, which she challenges.

picture is a further requirement that any individual have any particular beliefs or knowledge.¹¹ To get a feel for why Bird opts for such liberal position here,¹² we need only consider his case of Dr. N., which will be important to Lackey's critique.

Case of Dr. N. Dr. N. is working in mainstream science, but in a field that currently attracts only a little interest. He makes a discovery, writes it up and sends his paper to the *Journal of X-ology*, which publishes the paper after the normal peer-review process. A few years later, at time *t*, Dr. N. has died. All the referees of the paper for the journal and its editor have also died or forgotten all about the paper. The same is true of the small handful of people who read the paper when it appeared. A few years later yet, Professor O. is engaged in research that needs to draw on results in Dr. N.'s field. She carries out a search in the indexes and comes across Dr. N.'s discovery in the *Journal of X-ology*. She cites Dr. N.'s work in her own widely-read research and because of its importance to the new field, Dr. N.'s paper is now read and cited by many more scientists (Bird 2010: 32).

Bird's line on the case, in short, is that the scientific community knew the results of Dr. N's paper all along, despite there being a period of time where everyone who knew about—or indeed, was aware of—the result, was dead. Bird's diagnosis¹³ of the Case of Dr. N. clearly shows he takes the radical (SEK-R) to be true.

3. Lackey on (SEK-R) and Epistemic Defeat

Once Dr. N. and everyone else who read the paper dies, a lot can happen, some of it bad. Lackey exploits this fact with some complexity; consider her case *Addition*, which she takes to pose defeater-based problems that—as she argues—can't be overcome by (SEK-R), no matter what a proponent of (SEK-R) tries to do about it. Let *d* be a proposition representing the discovery Dr. N. is claimed to have

11. See Wray (2007: Section 4) for an argument against the claim that the scientific community, as a whole, can count as attaining collective knowledge, even if scientific research teams can possess collective knowledge.

12. Liberal, specifically, in the sense of what is epistemically tolerated of individuals in the group.

13. As Bird puts it, 'it seems irrelevant that Dr. N. and others who had read the original paper had died or forgotten about it. What is relevant is that the discovery was in the public domain, available, through the normal channels, to anyone, such as Professor O., who needed it' (Bird 2010: 32), cited also in (Lackey 2014: 284).

made, and which was subsequently forgotten.¹⁴ Now, here's Lackey's twist on the case:

Addition: Suppose that because of their ignorance of Dr. N.'s published paper, many members of the scientific community come to believe that not-*d* at T₂. Indeed, suppose that at scientific conferences and workshops, there is often explicit collective agreement among the participants that not-*d*. Because of this, the members of the scientific community act on not-*d* by, for instance, asserting that not-*d* in lectures and published work, approving cancer drugs that depend on not-*d*, conducting further experiments for cancer treatment that rely on not-*d*, applying for grants that take not-*d* for granted, and so on¹⁵ (Lackey 2014: 291).

Lackey takes it as obvious that, in *Addition*, the members of the scientific community believe that not-*d*; this is, she says, a "classic instance of a rebutting psychological defeater" (2014: 293)—which defeats the would-be knowledge and justification in virtue of being *believed*, and counting against the truth of the target proposition—*d*.¹⁶ With this observation in hand, she poses a dilemma to any would-be proponent of (SEK-R).

(SEK-R)	In <i>Addition</i> , is group knowledge psychologically defeated?
Yes	(i) Sceptical objection; (ii) Arbitrariness objection
No ₁	Group K(<i>p</i>) even though Group B(not- <i>p</i>)
No ₂	Deny belief is a necessary condition on knowledge

Obviously, given that *Addition* has all the signs of a case of rebutting psychological defeat, the first-horn of this dilemma looks *prima facie* best, but as Lackey argues, any would-be proponent of (SEK-R) who takes this horn in *Addition* will face two objections: a *sceptical* objection and an *arbitrariness* objection. The sceptical objection is, in short, that (SEK-R) can maintain the *prima facie* plausible result that the scientific community's belief that not-*d* is a psychological defeater for the alleged knowledge that *d* only by allowing that social knowledge can be *too easily* defeated, given the prevalence of dissenting opinion amongst scientists.¹⁷ How-

14. Perhaps Bird would object to the description of claim *d* as being 'forgotten' in this context. I am happy to replace this with 'no longer present in human memory'.

15. There is perhaps some scope for Bird to object that a case like this will not play out in practice. To the extent that Dr. N.'s published paper really is easily accessible by normal channels, then on the assumption that (as in the case of *Addition*) the scientific community really is very interested in matters connected to *d*, it's implausible (given actual scientific standards) that no one would have come across *d*. Thanks to Olle Blomberg for discussion on this point.

16. Psychological defeaters are also referred to by some writers under the description of *mental state* defeaters. See Sudduth (2008).

17. Of course, there is some room to press back here. For one thing, it might not be as implausible as it would first seem to simply allow that scientific *knowledge* is not as extensive

ever, even if this sceptical problem could be dealt with, another problem faces the proponent of (SEK-R) who takes the first horn: given that (SEK-R) implies that group knowledge needn't (as in the case of *Addition*) supervene on individuals' mental states, it would be epistemically arbitrary for mental states of individual members to contribute negatively but not positively to the group's mental state¹⁸.

The second horn, however, looks no more promising. This involves the proponent of (SEK-R) denying that knowledge is defeated in *Addition*. This route looks unpromising from the start, as anyone who takes it needs to explain why defeat *merely appears* to be present in *Addition*. Even if this can be done, it's unlikely it can be done without incurring further problems. Lackey here considers two routes: either deny that defeat is present at T_2 by submitting that a group can know that p even though the group believes that $\text{not-}p$; or deny that belief is a necessary condition on knowledge. As Lackey sees it, the first move here is in effect "exempting groups from the norms governing rationality and thus removing them from the realm of the rational altogether," (2014: 17) whereas the second preserves rationality only by failing to save any distinction between a group knowing and being in a position to know.

All in all, this is a pretty powerful argument Lackey has leveled against (SEK-R). Of particular interest in what follows is that where (SEK-R) is obviously lacking, (SEK-M) looks as though it would have no obvious problem handling epistemic defeat, given that a proponent of (SEK-M) looks poised to sidestep altogether the kind of defeater-objection raised by *Addition*-style cases. After all,

as we presume, while maintaining that this result would be compatible with insisting that *everyday* knowledge (of the sort that is of primary interest not to be ceded to the sceptic) is more or less what we thought. Lackey might rebut this point however by noting that, firstly, it's an unacceptable sceptical concession to cede *as much* knowledge to the sceptic as we would by allowing scientific knowledge to be defeated so easily. Secondly, it's plausible that there is some significant overlap between the scientific knowledge that would be ceded to the sceptic and everyday knowledge. Thanks to Duncan Pritchard for helpful discussion on this point.

18. A point of clarification. (SEK-R) is compatible with knowledge supervening on individuals' mental states. At least, it's not clear from Bird's description of what must be present in cases of socially extended knowing that supervenience on mental states of individuals would be excluded in principle. The crux of the point Lackey is highlighting is that, at least in some cases, group knowledge can be present even when no individual's mental state contributes to the group's positive epistemic status. And so the arbitrariness charge should be best understood as restricted to cases where mental states do not, in fact, contribute positively to group knowledge, though they contribute negatively. Furthermore, there is scope to argue, *contra* Lackey, that it's not true that individuals' mental states do not contribute *at all* to the group knowledge that by (SEK-R) is claimed to be present in the Case of Dr. N. After all, the relevant discovery that was submitted into the scientific oeuvre was due to the intellectual efforts of Dr. N. Therefore, at least one individual's mental states contribute (qua causal antecedent) to any group knowledge present later. Presumably, Lackey's arbitrariness charge must reflect the worry that individuals' mental states can contribute negatively to a group mental state that counts as knowledge at some time, t , even though, at t , no individual's mental states are contributing positively toward that state.

one who endorses (SEK-M) has space to claim that once there was no one around to believe the target proposition d , no knowledge of d was present.¹⁹ Accordingly, Lackey does *not* try to foist a dilemma on the proponent of (SEK-M) as she did with (SEK-R), as the dilemma is premised upon the acceptance of the presence of group knowledge that d . Thus Lackey does not take the defeat-based issues she's raised to make trouble for the more moderate position.

4. A New Argument from Defeat against (SEK-M)

I want to now suggest that it's not obvious at all just how the more modest (SEK-M) is going to be able to make sense of epistemic defeat any better than the more radical proposal. First, I'll advance an argument that (SEK), given certain background assumptions, can't make sense of psychological defeat. I'll then advance a similar argument with respect to normative defeat.

4.1. Psychological Defeat

The argument I want to advance for the claim that (SEK-M) can't adequately be squared with ordinary thinking about psychological defeat takes shape once two background claims are in place—claims which many collective epistemologists either explicitly or tacitly embrace. The first background claim is a very standard way to think about the kind of joint *belief* that is present in cases of group knowledge. Call this the *Joint Acceptance Model of Group Belief* (JAB).²⁰

Joint Acceptance-Belief (JAB): (i) A group, G , believes p iff the members of G jointly accept p ; (ii) the members of G jointly accept that p when the members conditionally commit to accept that p ; (iii) members of G conditionally commit to accept that p when each is committed to acting as if²¹ p provided the others do.²²

19. Recall that (SEK-M), while allowing that a group can know that p even when no individual in the group knows that p , can permit epistemic requirements weaker than knowledge on individuals—requirements that individuals clearly do not meet while dead

20. See here especially Gilbert (1987; 2002; 2013). For an alternative non-summativist account of group belief, see the positional account of group beliefs as defended by Tuomela (1995).

21. As Mathiesen writes, capturing this idea, "in contexts where her group membership is salient she is committed to act, speak, and reason consistently with the group view. If she fails to do so, then she has violated an obligation she has to her fellow group members." (2006: 196) Gilbert (2004: 104) gives an illustrative example: "The members of a seminar on human rights may in the course of a meeting form a joint commitment to believe as a body that the notion of a group right is a viable one. This would involve a requirement to express that belief at least within the confines of the seminar when it is in session. More broadly, it would require that each party express that belief when acting as a member of the seminar."

22. Note that one can act as if p without positively believing oneself, in a private capacity,

The second background claim is what Lackey herself endorses under the description of the *Group Member Action Principle*, a principle that says, in short, that a group can't do something if its members do nothing.

Group/Member Action Principle (GMAP): For every group, G , and act, ϕ , G performs ϕ only if at least one member of G performs some act or other that causally contributes to ϕ .²³

With (JAB) and (GMAP) in hand, there is a straightforward reductio-style argument to the effect that psychological defeaters cannot plausibly be accommodated on at least the most typical construals of (SEK-M), on which group knowledge involves group belief. The argument begins predictably enough.

- (1) Group A believes p only if A 's members are committed to acting as if p . [from (JAB)]
- (2) A can acquire the belief that q , where q indicates that not- p ,²⁴ only if at least one individual performs some action that causally contributes to A 's acquiring the belief that q . [from (GMAP)]
- (3) Any such action in (2) violates one's commitment to acting as if p .²⁵

Premises (1) and (2) are implied by (JAB) and (GMAP), respectively. Premise (3) makes explicit a further point that should be troubling for a proponent of the conjunction of (JAB) and (GMAP): that doing something that contributes to the group's acquisition of a belief that indicates not- p is an instance of not acting as if p . (It is rather acting, by contrast, as if not- p ²⁶).

that p ; believing not- p in a private capacity is not at conflict with joint commitment. As 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 baldly. (Gilbert 1989: 306-7)

23. Lackey notes that this position is embraced by, among others, List and Pettit (2011: 64).

24. I am using this language in a way that is meant to be inclusive between rebutting and undercutting psychological defeaters, though to make things simple, I'll be focusing on psychological defeaters of a rebutting variety.

25. What accepting p prohibits is acting, but acting in a way that is *not* as if p is true. Note that if one does nothing at all, one is not acting. I say this to clear up any confusion that the requirement that one act as if p is true implies the implausible result that one has to always be acting as if p is true in the sense that one must always be acting.

26. Consider an example in the individual case. Suppose I plan to place a bet on the Yankees-Braves game. I believe the Yankees will lose to the Braves and accordingly place a

Two quick points. Firstly, it should be plain to see that (JAB) and (GMAP), paired together, render it confusing how to make sense of inconsistent group beliefs. I've argued for this point elsewhere²⁷ but won't rehearse these arguments here. Secondly, notice that there is no analogous problem in the case of non-collective (individual) epistemology. Believing that p but acting as if not- p is an instance of irrational action, where irrational action in accordance with not- p is, on orthodox thinking in individual epistemology, compatible with the original belief that p .²⁸

That said, note that (1-3) imply:

- (4) Therefore, if A acquires the belief that q , where q indicates that not- p , it's not the case that A retains its group belief that p .

With premise (4) in hand, it's then a straight shot to the conclusion that a proponent of (SEK-M) who signs on to (JAB) and (GMAP) can't account for psychological defeat. Here's the rest of the argument.

- (5) Group A 's belief that q , where q indicates that not- p , can function as a psychological defeater for A 's belief that p only if A acquires the belief that q while retaining A 's belief that p .²⁹
- (6) Therefore, A 's belief that q cannot be a psychological defeater for A 's (original) belief that p . [from 4,5]

It looks like (SEK-M) can make sense of psychological defeaters then only by rejecting either the very plausible Group/Member Action Principle (GMAP) or by rejecting the standard account of group belief, (JAB).³⁰ Many working in collective epistemology will not be happy with either of these options.

bet on the Braves. I then acquire a new belief to the effect that the entire Braves team has just acquired malaria, and no player can currently stand up, and act accordingly, by placing my bet on the Yankees. In doing so, I am no longer acting as if the Braves will win.

27. See Carter (2014).

28. Thanks to an anonymous referee at *Ergo* for comments that led to an improvement of this presentation.

29. To see why this is so, just consider that if I believe p at t_1 , but stop believing p at t_2 , and then acquire q which indicates not- p at t_3 , my belief that q is not a psychological defeater for my belief that p , given that I no longer have a belief that p , at t_3 , when I acquire q .

30. It is worth noting and replying to a potential line of objection. Suppose that S is part of an NGO that is trying to fight groups that argue that CO₂ doesn't contribute to rising global temperature. Stipulate that S 's group has knowledge that CO₂ does contribute to rising global temperature. But as a consequence of S 's efforts to learn more about the arguments of the opponent groups (in order to rebut them), S acquires a belief that undermines some of the key data points that supported the group belief that CO₂ emissions drive a rise in global temperature. S takes this up for discussion in S 's group which collectively accepts that those data points are unreliable, and S 's group agrees not to rely on them in future. One might argue that this could be a psychological defeater of the group's (previous) knowledge that CO₂ emissions cause rising global temperature, but yet, it is plausible that the group

What are the escape routes? One is to reject (GMAP). That's not very plausible. What about rejecting (JAB)? Here there is some wiggle room. Rejecting (JAB) is precisely what is done by rejectionists about group belief, who typically advert to alleged ways that belief and acceptance come apart at the individual level to insist that what is present in the case of groups is more akin to acceptance than belief.³¹ Rejectionism has a negative component and a positive component:

Rejectionism: For any group, G , and proposition, p : (i) G cannot believe that p ; (ii) G can accept that p .³²

What does (SEK-M) look like when paired with rejectionism? A simple way to think about the pairing is put forward by Hakli (2007), who argues that group knowledge is to be defined as justified true *acceptance*.³³ A Hakli-style version of (SEK-M) then says that a group, G , can know that p even when not a single individual member of G knows that p , but *not* if no individual member of G *accepts* that p .

Now, can a 'Hakli-style' rejectionist version of (SEK-M) make sense of psychological defeat better than a proponent of (SEK-M) who embraces (JAB)?³⁴ Obviously, if psychological defeat is *essentially* defined in terms of belief, then (SEK-M) couldn't resort to rejectionism (about belief!) in order to make sense of it. But let's grant (charitably) that the rejectionist proponent of (SEK-M) has a card to play: let's say the rejectionist proponent of (SEK-M) counts as making

remains collectively committed to acting as if it is true that CO₂ emissions cause rising global temperature; the group simply takes it that not accepting the data points is compatible with this commitment. In response to this line, it is important to note that the devil is in the details of the case. If the data points in question are but a few data points among hundreds on which the NGO bases its collective belief that CO₂ contributes to rising temperatures, then not accepting the data points is indeed compatible with the commitment to act as if CO₂ emissions cause rising global temperature. Though in such a case it is implausible that the knowledge in question actually is defeated. Compare: in the individual case, if I base my belief that Team A will win on 100 items of information, and come to doubt 3, my overwhelming evidence that Team A will win remains. If, however, (in the NGO example) the data points in question were the three *central* data points on which the group based its belief that CO₂ emissions are present, then it is more plausible that knowledge would be defeated. However, in this variant of the case, it is no longer plausible to suppose that in acquiring the group belief that the data points are false, the group continues to retain its original belief that CO₂ contributes to rising temperatures. Thank you to Olle Blomberg for raising this objection.

31. For a discussion of differences between accepting and believing, at the individual level, see, for example, Engel (1998). Cf. Cohen (1992).

32. Proponents include Meijers (2002), Wray (2001) and Hakli (2007).

33. See also Wray (2007) for a defence of the view that research teams are candidates for (non-summativist) group knowledge understood as justified true acceptance. Wray however denies that such candidate collective knowers include communities of scientists identified with research fields.

34. Note that such a version of (SEK-M) is not committed to (JAB), which, in conjunction with (GMAP), was part of the original problem for (SEK-M).

sense of psychological defeat provided she can reconcile group knowledge (on the rejectionist version of (SEK-M)) with *quasi-psychological defeat*, where the idea in play is that knowledge that p can be quasi-psychologically defeated by S 's acceptance (rather than *belief*) that q , where q indicates not- p .

But unfortunately this attempt to save (SEK-M) by retreating to rejectionism runs quickly to a dead end. We can simply run a variation of the previous argument to establish that the proponent of (SEK-M) who embraces a rejectionist line can't make sense of quasi-psychological defeat, as we've just defined it. The argument, modeled upon the previous argument, runs as follows:

- (7) Group A accepts p only if A 's members are committed to acting as if p . [rejectionist-friendly version of (1)]
- (8) A can acquire the acceptance that q , where q indicates that not- p , only if at least one individual performs some action that causally contributes to A 's acquiring the acceptance that q . [from (GMAP)]
- (9) Any such action in (8) violates one's commitment to acting as if p .
- (10) Therefore, if A acquires the acceptance that q , where q indicates that not- p , it's not the case that A retains its group acceptance that p .
- (11) Group A 's acceptance that q , where q indicates that not- p , can function as a (quasi)-psychological defeater for A 's acceptance that p only if A acquires the acceptance that q while retaining A 's acceptance that p .
- (12) Therefore, A 's acceptance that q cannot be a (quasi)-psychological defeater for A 's (original) acceptance that p . [from 10,11]

4.2. Normative Defeat

At this point, (SEK-M) isn't faring much better than (SEK-R) did, in so far as its prospects are looking for making sense of psychological defeat.

What about *normative* defeat? (Maybe one out of two isn't so bad?) Following Lackey, let's define a normative defeater as follows:

Normative Defeater: A doubt or belief that S ought to have, and that indicates that S 's belief that p is either false (i.e., rebutting) or unreliably formed or sustained (i.e., undercutting).

Normative defeaters defeat knowledge and justification by virtue of being doubts or beliefs that S *should* have (whether or not S does have them) given the presence of certain available evidence.³⁵ Important to the matter of whether knowledge is

³⁵ Lackey discusses normative defeat in more detail in (among other places) Lackey (1999; 2005).

vanquished by normative defeaters will be the nature of the ought that features in normative defeat. We'll get to that. But first, a quick example of a normative defeater. Suppose you are on your way to the zoo and—through inattention and carelessness—overlook several signs, each of which says the zoo is closed. Even if you had excellent evidence previously that the zoo would be open, your epistemic justification (and/or would-be knowledge) that the zoo is open is defeated in this case by a normative rebutting defeater, a belief you *ought* to have had, and which counts against the truth of the target proposition. In the undercutting case, just suppose you've begun taking a certain experimental medicine of which a potential side-effect (clearly stated on the bottle) is Capgras delusions. You don't read about this side effect, and begin to convince yourself that your family members have been replaced by identical imposters. In this case, whatever positive epistemic status your belief might aspire to is undercut by a normative defeater: you ought to have been aware of the side effects, which (were you aware) would have indicated that your beliefs were not to be trusted. Plausibly, a group normative defeater will have to be a belief a *group* ought to have, and which defeats group beliefs via the same mechanisms as in the individual case. Define a group normative defeater as follows:

Group Normative Defeater (GND): A (group) normative defeater for group *A*'s belief that *p* is some belief that *A* ought to have, that *q*, where *q* indicates that *A*'s belief that *p* is *either* (i) false (i.e., rebutting) or (ii) unreliably formed or sustained (i.e., undercutting).

Let's now consider a reasonably straightforward argument that (SEK-M) can't account for normative defeat. The target of the argument, to be clear, will be any proponent of (SEK-M) who embraces, along with (GMAP) and (JAB), a further key background claim, which has to do with the kind of ought that features in normative defeat. What is of relevance here is the clear sense in which practical goals and desires seem entirely orthogonal to the ought of normative defeat, *qua epistemic ought*.

To make this point more concrete, consider here Feldman's (2000) remarks on the nature of the ought that features in the evidentialist's thesis that one ought to believe in accordance with one's evidence:

If the oughts in question are supposed to be means to goals that people actually have, then it seems that only people who do have the epistemic goals just mentioned would be subject to the relevant epistemic requirements. However [the claim that one ought to believe in accordance with one's evidence] is not restricted in that way. It says that all people epistemically ought to follow their evidence, not just those who have adopted some specifically epistemic goals. (2000: 682)

Feldman's observation about the categoricity of the *epistemic* ought goes hand in hand with ordinary thinking about epistemic (as opposed to practical) reasons.³⁶ Here is Railton (1997):

On the usual view of things, two agents in the same epistemic situation (same evidence, same background beliefs) would have the same reason for believing any given proposition, regardless of possible differences in their personal goals.³⁷ (1997: 53)

The relationship between epistemic oughts and epistemic reasons is one of considerable disagreement amongst value theorists,³⁸ though it needn't bog us down here. The crux of the foregoing can be expressed simply: the ought of normative defeat is *categorical*—it does not depend on our having any antecedent practical goals or commitments. And this simply follows from two more basic ideas: that the ought of normative defeat is an epistemic ought, and secondly, that epistemic oughts are categorical, not hypothetical.

We are now in a position to state the argument that (SEK-M) can't make sense of normative defeaters. Assume, for *reductio*, that the proponent of (SEK-M) says the following:

Supposition: For some Group *G*, which believes *p*, *G*'s belief that *p* is, *qua* would-be knowledge, *normatively* defeated.

From this supposition, it would follow that there is a *q* such that group *G* (categorically) ought to believe *q*, where the truth of *q* counts against the truth of *p* or indicates *p* was unreliably formed or sustained. This, after all, follows from (i) the definition of (group) normative defeat (GND), and (ii) the observation that the ought of normative defeat is the epistemic ought.

Now, if we take seriously the idea that there is a *q* such that group *G* (categorically) ought to believe *q*, where the truth of *q* counts in some way against the truth of *p*, then it follows from (GMAP) that at least one member of *G*, *M* ought to do some action, *φ*, that causally contributes to bringing it about that *G* believes that *q*, where *φ* involves acting as if *q*.

But a dilemma materialises once we consider the status of the 'ought' that features in the claim that *M* ought to *φ*:

In short, say the ought in question is *non-categorical*, and it looks like you're no longer referring to the ought of normative defeat. Of course, what's left is to go the other route and to embrace the categoricity of the ought in question. But

36. For more on this distinction, see Raz (2009) and Olson (2009).

37. Cf. Cuneo (2007): "if reason *R* is an epistemic reason to believe that *p*, for *S*, then this is so regardless of whether *S* desires to believe the truth."

38. For a recent and very comprehensive discussion of this relationship, see Skorupski (2010).

	What is the status of 'M ought to ϕ '?
Non-categorical	Not the ought of normative defeat
Categorical	Categorical and non-categorical

then a different problem takes shape. Given (JAB), any commitment M would have, qua member of the group G , to acting as if q (e.g. any commitment to ϕ) will always be *conditional*, in the sense, that M will, qua member of G , be only ever be committed to acting as if q provided others do.³⁹ But then it seems that the claim that M ought to ϕ is such that it both (i) must be categorical, and (ii) must not be categorical. Thus, the supposition that normative defeat occurs will, on the package that includes (SEK-M) paired with (JAB) and (GMAP) lead to consequences that cannot be reconciled with one another.⁴⁰

Let's consider now a potential objection to the reductio just offered. Call this objection the 'different targets' objection. Here's the worry, put abstractly: the principle according to which epistemic oughts are categorical applies, in the case of group belief, just to the group as a whole. However, the conditional part of the joint belief principle applies to individuals. But then, as this line of thought goes, there can be no inconsistency in taking the second horn of the dilemma posed above. This is because there is no inconsistency in the group being categorically committed to believing something an individual in the group is merely conditionally committed to believing.⁴¹

In response to this envisioned reply, it is important to note that a group categorically ought to do something, ϕ , only if there is some action, ψ , such that there is a member of group, M , such that M , categorically ought to do ψ , where M 's doing ψ causally contributes to the group's doing ϕ . This is an implication of embracing (GMAP). And it's this individual's action, ψ , which ends up with the strange status of being something that this individual both categorically and non-categorically ought to do, on the assumption that there is a normative defeater for the group and we take the second horn of the dilemma.

Once this point is appreciated, the different targets objection disappears. However, one might attempt to press things further. If there is something such

39. Compare here Gilbert's discussion of the obligation Jack has when jointly committing to walking with Jill. When Jack does something that's not compatible with walking together—for instance, walking faster than will be reasonably possible for Jill to walk along side him—then Jill can rightly rebuke him. Plausibly, Jill then is released from any obligation she had to walk at a pace accommodating Jack. See Roth (2011: Section 4)

40. The kinds of ϕ s I have in mind here will just be actions such that acting as if q involves these actions. Of course, some commitments one has, qua member of the group, can be commitments to do things other than *actions*, proper—e.g., commitments to be certain ways. This though is compatible with the point I'm making here, which relies just on the (GMAP) insight that the *group* can't do anything without some individual acting.

41. Thanks to an anonymous referee at *Ergo* for pressing this line of reply.

that the group ought to believe that thing, might it just be that *all* of the members are categorically committed to bringing it about that the group believe the thing in question, and not merely non-categorically committed to bringing this about?

There's really no room for a proponent of (JAB) to accommodate such an idea. Consider, for example as Gilbert (1993) puts it, that:

A joint commitment is "joint" in a strong sense: no individual is committed until all the others are; it is impossible to rescind the commitment to one party without rescinding it for the others; if the commitment is broken, it is broken for all: it does not remain to require any action of anyone. (1993: 693)

The logic of (JAB) is one that excludes an individual being non-conditionally committed to something as a member of a group. And so taking the second horn of the dilemma does ultimately generate the precarious consequence that at least one individual in the group have some obligation that is both categorical and non-categorical.

Let's consider another escape route. The reductio just considered affects any version of (SEK-M) that embraces (JAB) and (GMAP), as well as the categoricity of the ought of normative defeat. Again, (JAB) was part of the problem. Question: is there any wiggle room for a Hakli-style rejectionist version of (SEK-M) to preserve the insight that group knowledge should be subject to normative defeat? Recall, the 'Hakli' package includes (SEK-M), (GMAP), and a rejection of (JAB). Group knowledge = justified true (group) *acceptance*. Again, given that (GND) is framed *explicitly* in terms of group *beliefs* a group ought to have, Hakli could object to the legitimacy of having to account for it, as stated.

Presumably, though, if a Hakli-style version of (SEK-M) is going to have any chance of accommodating normative defeaters, it must be able to accommodate something like the following:

(Group) Normative Defeater-(Rejectionism) (GND-R): A (group) normative defeater for Group *A*'s *acceptance* that *p* is some *acceptance* that *A* ought to have that *q*, where *q* indicates that *A*'s acceptance that *p* is *either* (i) false (i.e., rebutting) or (ii) unreliably formed or sustained (i.e., undercutting).

The problem is, the very same structural problems that prevented an account of (SEK-M) on which group *beliefs* are required for knowledge from making sense of (GND) crop up for the proponent of a rejectionist Hakli-style version of (SEK-M) that tries to preserve that group knowledge can be defeated by (rejectionist-friendly) normative defeat, as specified in (GND-R). We can simply run a variant on the previous version of the reductio—one which will culminate, *mutatis mutandis*, in the same dilemma that is faced by a proponent of (SEK-M) who embraces

(JAB).⁴²

The upshot here is that it looks like the phenomenon of epistemic defeat (in both its psychological and normative varieties) renders the more plausible (SEK-M) in trouble no less than the more radical (SEK-R). This is bad news for orthodoxy in collective epistemology. It makes it harder to swallow that it's really *knowledge* collective epistemologists are talking about.

Having set up this problem, I now want to change course. There is, I think, yet a further escape route remaining, though it requires at least some departure from ordinary thinking about what group belief involves.

5. A New Manoeuvre

I conclude in this section by first sketching some key features of a model of group belief and knowledge I've defended elsewhere,⁴³ and then showing how this model (in virtue of these features) offers some *prima facie* promise for a proponent of (SEK-M) who seeks reconciliation with ordinary thinking about defeaters.

The shape that the model I propose takes is best presented as a kind of response to an oversimplification betrayed by one very natural strand of rejectionist thinking. For a clear sense of this strand, consider Preyer (2003: p. 9): "Beliefs are gradual (quantitative) and have subjective probabilities, acceptance is a matter of quality . . . They have a feature of all or nothing." Reiterating this latter observation, rejectionist Meijers (1999: p. 81) puts the point tersely: "We agree, or we do not."⁴⁴

A simple line of argument falls out of these rejectionist insights: we should

42. To see how this would go, simply assume, for reductio, that for some group *G*, which accepts that *p*, *G*'s acceptance that *p* is, qua would-be knowledge, *normatively* defeated. If this were so, it would follow that there is a *q* such that Group *G* (categorically) ought to accept *q*, where the truth of *q* counts against the truth of *p* or indicates *p* was unreliably formed or sustained. This much follows simply from (i) the definition of (group) normative defeat (GND-R), and (ii) the observation that the ought of normative defeat is the epistemic ought. Now, if we take seriously the idea that there is a *q* such that group *G* (categorically) ought to accept *q*, where the truth of *q* counts in some way against the truth of *p*, then it follows from (GMAP) that at least one member (*M*) ought to do some action, *φ*, that causally contributes to bringing it about that *G* accepts that *q*, where *φ* involves acting as if *q*. And now we've reached a variant of the same dilemma that emerges for the proponent of (SEK-R) who embraces (JAB). That is, what is the status of '*M* ought to *φ*'? If the ought in play is non-categorical, it is not the ought of normative defeat. If it is categorical, then this seems irreconcilable with the fact that any commitment *M* would have, qua member of the group *G*, to acting as if *q* (e.g., any commitment to *φ*) will always be *conditional*, in the sense, that *M* will, qua member of *G*, be only ever be committed to acting as if *q* provided others do.

43. See Carter (2014).

44. See also Meijers (2003) for some further defence of this rationale for rejectionism, in light of challenges from Gilbert (2002) and others.

reject that groups have beliefs, because whatever groups accept, *qua* agreements, they accept in an all-or-nothing fashion, where as beliefs don't have such an all-or-nothing feature; rather, they are gradient.⁴⁵

Though there is a clear sense in which agreement is (as Meijers suggests) an on-or-off affair, what is easily overlooked—by rejectionists and (JAB) proponents alike—is that there is significant flexibility with respect to the *object of agreement*, when that object is a representational attitude. And this is because there are very different ways we can rely on *p*, ways that themselves correspond with different levels of practical risk.

To unpack this point, consider a rather mundane observation in individualist epistemology concerning the relationship between credence, or degree of belief, and rational action. Contrast two cases where Stan believes his wallet is in his pocket. In the first case he is *certain* that his wallet is in his coat pocket (call this *p*), and in the second case, he is about .7 confident his wallet is in his coat pocket.⁴⁶ Stan's epistemic position in the first case rationalises Stan's acting as if *p* under a *wider* range of practical circumstances than does the second, where his epistemic position is not taken (by Stan) to be as good. Put differently, Stan's being certain that *p* rationalises acting as if *p* across situations of *higher practical risk* than will be rational for Stan to continue to act as if *p* when Stan's epistemic position is regarded as lower, as in the second case. Accordingly, if it's very important to be able to make a payment, Stan might well be rational in acting as if *p* if his situation is as it is in the first case, even if not in the second (even though Stan counts nonetheless as *believing* his wallet is in his pocket in the second case).⁴⁷

This insight has interesting and fruitful applications in the arena of group belief.⁴⁸ Here is the central statement of the view I've defended elsewhere: A group can believe *p* in virtue of jointly accepting that *p*, where this involves an agreement to act as if *p* in the following qualified way: to act as if *p* across a *range of practical circumstances that positively correlates with the group's epistemic position*, which is a position the individuals of the group (in jointly accepting that

45. And there will be some further premise here to the effect that gradience is the kind of property beliefs have *essentially*, if at all.

46. I am assuming that Stan counts as believing the proposition when his credence is .7. While this assumption is certainly friendly to a simple Lockean threshold view of belief, where one counts as having a belief once one's credence is sufficiently high, it's not required in order to illustrate the point I am making here about the relationship between rational credence and rational action. All that point requires is that we assume he counts as believing in both cases, and that this is at least compatible with altering his credence from 1 to .7.

47. For a detailed discussion of such cases, and more generally of the connection between rational belief and rational action, see Carter et al. (in press).

48. To draw a parallel from the case of Stan, consider two extremes: in the first case, (and at one limit) we might jointly agree as a group to endorse *p without qualification* regarding just *how* we shall act as if *p*. At the other limit, we might agree to jointly accept (and thus act as if, and rely on) *p*, only if there is little to no practical risk to being wrong.

p) are obligated, qua members of the group, to update in systematised ways.

Call this the *risk responsiveness account of group belief*. This is not the place to provide a full defence of the view,⁴⁹ though I want to conclude by gesturing toward some novel resources this proposal has for the purposes of making sense of epistemic defeat, where the previously considered approaches seemed lacking.

In particular, I'll conclude by sketching how the risk-responsiveness approach to group belief stands to, at least potentially, when paired with (SEK-M), make sense of group psychological defeaters. The discussion here should make it clear how similar moves can be made to make sense of normative defeat, though I won't argue for this here.

We can envision a move on the table that proceeds as follows: group psychological defeaters function by virtue of being had, and counting (in some way) against the *group's continuing to hold belief in the way it is held*, and regardless of the truth-value or epistemic status of the defeater. On this picture, *individual updates* can (on the risk-responsiveness model) count against holding belief in the way it is held, and so can defeat group knowledge in this sense, and regardless of epistemic status or veracity of these updates.

This point can be illustrated by an example:

Scholarship Committee: A scholarship committee, SC, is going to distribute three scholarships: the gold scholarship (£100,000), the silver scholarship (£50,000) and the bronze scholarship (£25,000) to three finalists: students *A, B* and *C*. The decision will be made on the basis of an estimated weighted average of financial need (50%) and academic potential (50%). In order to divide labour, some members of the committee researched the dossier of *A*, others *B* and others *C*. Information is pooled into a central database, and the group does two things. Firstly, they systematically analyse all pooled information. Secondly, they vote on how best to characterise their epistemic position (on the basis of reliability and comprehensiveness of information, depth of analysis, etc.) with respect to the matter of which student merits which award. They accordingly rate their position a 7/10. Finally, the committee embraces a positive view (on the basis of their available information and analysis) as to which candidate, in light of the stated rules, merits which scholarship. Suppose that after pooling resources and weighing information (in a structured way), the group agrees to jointly accept the following position: that *A* merits the gold, *B* merits the bronze and *C* merits the silver.

Now, let us add two variations on this case.

Variation-1: Suppose Rebus is a member of SC, and that Rebus overhears some of the committee members discussing how the candidate *B*'s ACT

49. For a more detailed discussion of this kind of view, see Carter (2014).

test scores were actually much higher than they reported—a perfect 36—enough (from what Rebus can tell) to potentially move B past candidate A for the gold scholarship. Deeply concerned, Rebus independently confirms this and accordingly updates the committee’s database to reflect candidate B’s ACT score was in fact a 36 (rather than the 27 that was reported). In independently confirming this, Rebus also uncovers (and records) that candidate B has a history of being a successful plaintiff in a high-profile discrimination suit.

Variation-2: Rebus finds out that a committee member, Uriel (in charge of collecting data on candidate A) is disgruntled and unstable. Uriel reveals verbally that, as this is his final year on the scholarship committee, he is ‘phoning it in’, but gives no further details. Rebus updates the committee’s website noting that there is reason to believe that the individual responsible for collecting information on candidate A was unreliable.

I think a case can be made, with reference to the risk-responsive model of group belief, that *Variation-1* is structurally very closely akin to a *rebutting psychological defeater*, and likewise *Variation-2* an *undercutting psychological defeater*, for the group belief that the three scholarships merit award to A, B, and C (in that order). And, furthermore, that this result is achieved in a way that avoids the kinds of objections canvassed in the previous sections to more traditional proposals embracing (SEK-M).

Consider that, in *Variation-1* the updated information, which Rebus qua member of the group) made available to the group (via the central database), played two roles, to update that the group’s epistemic position *vis-a-vis* the proposition that the three scholarships merit award to A, B, and C (in that order) is not as good as previously thought (previously, .7), and secondly, that the practical risk in play is in fact greater than previously anticipated. These updates, combined, very plausibly count against the group’s continuing to rely on the group belief that the scholarships merit award in the order previously accepted, across the range of practical circumstances that they would have previously relied on it. A proponent of (SEK-M) who insisted that, previously, group knowledge was present in the case now has a story to tell for how that knowledge is defeated. To the extent that—as is a popular position—knowledge is *actionable*⁵⁰ (given present levels of practical risk), the group’s information state (following Rebus’s updates) is not knowledge.

One important benefit of this story is that by replacing (JAB) with the risk-responsiveness account of belief, a certain problematic consequence can be avoided without rejecting (GMAP). To appreciate this point, suppose one were to attempt

50. See Hawthorne and Stanley (2008) for the most widely discussed version of this position.

to redeploy the argument from Section 4.1 against the risk-responsiveness account. Here's how the redeployed argument would go:⁵¹

- (13) Group *A* believes *p* only if *A*'s members are committed to acting as if *p* across a range of practical circumstances that positively correlates with the group's epistemic position, which is a position that the individuals of the group (in jointly accepting that *p*) are obligated, qua members of the group, to update in systematized ways [From risk responsiveness account].
- (14) *A* can acquire the belief that *q*, where *q* indicates that not-*p* only if at least one individual performs some action that causally contributes to *A*'s acquiring the belief that *q*. [From GMAP]
- (15) Any such action in (14) violates one's commitment to acting as if *p*.
- (16) Therefore, if *A* acquires the belief that *q*, where *q*, indicates that not-*p*, it's not the case that *A* retains its group belief that *p*. [from 13–15]
- (17) Group *A*'s belief that *q*, where *q* indicates that not-*p*, can function as a psychological defeater for *A*'s belief that *p* only if *A* acquires the belief that *q* while retaining *A*'s belief that *p*. [by definition]
- (18) Therefore, *A*'s belief that *q* cannot be a psychological defeater for *A*'s (original) belief that *p*. [from 16–17]

If we were working with (JAB) rather than the risk-responsiveness account, (15) would come out true (see Section 4.1). However, on the risk-responsiveness account, (15) is false. On the risk-responsiveness model, actions that play the role of causally contributing to the group's acquiring the belief that *q* can be actions individuals within the group are committed to, qua group members, that believe the target proposition (*p*). For example, part of what it is to be a member of a group that believes that *p* on the proposed account is to update the group's epistemic position in connection with the practical circumstances in the manner suggested. And as I've suggested, such updates can play the role of psychological defeaters for the group's continuing to believe the target proposition by acting as if it's true *in the way* they had before.

What about *Variation-2*? With reference to the risk-responsive model of group belief, *Variation-2* is structurally very closely akin to a *undercutting psychological defeater* for the group belief that the three scholarships merit award to A, B, and C (in that order). The rationale proceeds as follows. In *Variation-2* the updated information, which Rebus qua member of the group) made available to

51. Thanks to an anonymous referee at *Ergo* for suggesting this line be made explicit.

the group (via the central website) regarding Uriel's unreliability very plausibly counts against the reliability of the mechanisms giving rise to the group's relying on the group belief that the scholarships merit award in the order previously accepted, across the range of practical circumstances that they would have previously relied on it. In short, then, a proponent of (SEK-M) who insisted that, previously, group knowledge was present in the case, now has a story to tell for how that knowledge is defeated. Again, to the extent that knowledge is *actionable* (given present levels of practical risk), the group's information state (following Rebus's updates about Uriel) is not knowledge.⁵²

6. Concluding Remarks

In conclusion, I hope to have illuminated a worrying, and thus far overlooked, weak spot which collective epistemologists have yet to adequately address. I've taken as a starting point a suggestion intimated by Lackey, which is that an *ex ante* constraint on any plausible account of group knowledge is that it be in principle defeated through ordinary mechanisms of epistemic defeat. Whereas Lackey has shown how Bird's radical construal of socially extended knowledge fails this constraint, I've argued here that, given assumptions widely embraced in collective epistemology, the more moderate construal of the thesis faces just as much trouble. Though the central contribution here has been to raise this novel problem for the more moderate view, I've concluded by attempting to sketch a way forward. The way to accommodate epistemic defeat outlined in Section 6, granted, involves a number of commitments that proponents of group knowledge might not be amenable to. To the extent that they are not, it is incumbent upon proponents of group knowledge to find some alternative way to make sense of epistemic defeat. And to the extent that group knowledge is to be taken seriously, making sense of epistemic defeat is not a task that can be set aside.

52. The reason the risk-responsiveness approach offers theoretical possibilities not available to the more traditional (JAB) model is structural. The same structural move made to make sense of psychological defeat can be exploited in a way that has promising resources for making sense, as well, of normative defeat. The most helpful way to think about this is to highlight, on the risk responsiveness model, that there are very different ways a group can act as if *p*. While individuals are conditionally committed to acting as if *p* in a qualified way (in a way that aligns with the circumstances of risk), the group itself is best understood as categorically committed to *p*, simpliciter. In this respect, the flexibility of the risk responsiveness approach prevents an individual within the group to be both categorically and non-categorically committed to the same action.

Acknowledgments

Thanks to Olle Blomberg, Emma C. Gordon, Jesper Kallestrup, Jeroen de Ridder, Orestis Palermos, and Duncan Pritchard for helpful comments on earlier drafts. Thanks also to two anonymous referees at *Ergo* 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.

References

- Bergmann, Michael (1997). Internalism, Externalism and the No-Defeater Condition. *Synthese*, 110(3), 399–417.
- Bird, Alexander (2010). Social Knowing: The Social Sense of 'Scientific Knowledge'. *Philosophical Perspectives*, 24(1), 23–56.
- Carter, J. Adam (2014). Group Peer Disagreement. *Ratio*. Advance online publication. doi: 10.1111/rati.12077.
- Carter, J. Adam, Benjamin Jarvis, and Katherine Rubin (in press). Belief without Credence. *Synthese*.
- Cohen, L. Jonathan (1992). *An Essay on Acceptance and Belief*. Clarendon Press.
- Cuneo, Terence (2007). *The Normative Web: An Argument for Moral Realism*. Clarendon Press.
- Engel, Pascal (1998). Believing, Holding True, and Accepting. *Philosophical Explorations*, 1(2), 140–151.
- Feldman, Richard (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60(3), (667–695).
- Gilbert, Margaret (1987). Modelling Collective Belief. *Synthese*, 73(1), 185–204.
- Gilbert, Margaret (1989). *On Social Facts*. Routledge.
- Gilbert, Margaret (1993). Agreements, Coercion, and Obligation. *Ethics*, 103(4), 679–706.
- Gilbert, Margaret (2002). Belief and Acceptance as Features of Groups. *Protosociology*, 16, 35–69.
- Gilbert, Margaret (2004). Collective Epistemology. *Episteme*, 1(2), 95–107.
- Gilbert, Margaret (2013). *Joint Commitment: How We Make the Social World*. Oxford University Press.
- Hakli, Raul (2007). On the Possibility of Group Knowledge without Belief. *Social Epistemology*, 21(3), 249–266.
- Hawthorne, John and Jason Stanley (2008). Knowledge and Action. *The Journal of Philosophy*, 105(10), (571–590).
- Hutchins, Edwin (1995). *Cognition in the Wild*. MIT Press.
- Ichikawa, Jonathan Jenkins and Matthias Steup (2014). The Analysis of Knowl-

- edge. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014 ed.).
- Knorr Cetina, Karin (2009). *Epistemic cultures: How the Sciences Make Knowledge*. Harvard University Press.
- Lackey, Jennifer (1999). Testimonial Knowledge and Transmission. *The Philosophical Quarterly*, 49(197), 471–490.
- Lackey, Jennifer (2005). Memory as a Generative Epistemic Source. *Philosophy and Phenomenological Research*, 70(3), 636–658.
- Lackey, Jennifer (2012). Group Knowledge Attributions. In Jessica Brown and Mikkel Gerken (Eds.), *New Essays on Knowledge Ascriptions* (243–269). Oxford University Press.
- Lackey, J. (2014). Socially Extended Knowledge. *Philosophical Issues*, 24(1), 282–298.
- Lackey, J. (Forthcoming). *The Epistemology of Groups*. Oxford University Press.
- List, Christian and Philip Pettit (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- Mathiesen, Kay (2006). The Epistemic Features of Group Belief. *Episteme*, 2(3), 161–175.
- Meijers, Anthonie (1999). Believing and Accepting as a Group. In Anthonie Meijers (Ed.), *Belief, Cognition and the Will* (59–73). Tilburg University Press.
- Meijers, Anthonie (2002). Collective Agents and Cognitive Attitudes. *Protosociology*, 16, 70–80.
- Meijers, Anthonie (2003). Why Accept Collective Beliefs? Reply to Gilbert. *Protosociology*, 18–19(IV), 377–388.
- Olson, Jonas (2009). Error Theory and Reasons for Belief. In Andrew Reisner and Asbjørn Steglich-Petersen (Eds.), *Reasons for Belief* (75–93). Cambridge University Press.
- Pollock, John (1986). *Contemporary Theories of Knowledge*. Rowman and Littlefield.
- Preyer, Gerhard (2003). What is Wrong with Rejectionists? Retrieved from http://www.fb03.uni-frankfurt.de/48846514/GPreyer_Rejectionists1.pdf
- Railton, Peter (1997). On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action. In Garrett Cullity and Berys Gaut (Eds.), *Ethics and Practical Reason*. Cambridge University Press.
- Raz, Joseph (2009). Reasons: Practical and Adaptive. In David Sobel and Steven Wall (Eds.), *Reasons for Action* (37–57). Cambridge University Press.
- Roth, Abraham Sesshu (2011). Shared Agency. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2011 ed.).
- Skorupski, John (2010). *The Domain of Reasons*. Oxford University Press.
- Sudduth, Michael (2008). Defeaters in Epistemology. *Internet Encyclopedia of Philosophy*.
- Tollefsen, Deborah Perron (2006). From Extended Mind to Collective Mind. *Cog-*

nitive Systems Research, 7(2), 140–150.

Tuomela, Raimo (1995). *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press.

Wray, K. Brad (2001). Collective Belief and Acceptance. *Synthese*, 129(3), 319–333.

Wray, K. Brad (2007). Who Has Scientific Knowledge? *Social Epistemology*, 21(3), 337–347.