

Collective Epistemic Responsibility: a Preventionist Account

Will Fleisher¹ and Dunja Šešelja²

¹*Washington University, St. Louis, fleishwp@gmail.com*

²*Philosophy and Ethics Group, TU Eindhoven, d.seselja@tue.nl*

March 13, 2020

Abstract

If a given scientific community faces an epistemic harm that could be prevented only by a collective action, what kind of epistemic duties fall on each of the given scientists? In this paper we propose an account of collective epistemic responsibility, which addresses this and related questions. Building on Hindriks' (2018) account of collective moral responsibility, we introduce the Epistemic Duty to Join Forces as a norm consisting of two sub-norms: first, a duty of individuals to approach other relevant agents raising awareness about the epistemic harm, expressing willingness to prevent it, and encouraging others to do the same; and second, a duty of those who have expressed their commitment to join forces, to prevent the given epistemic harm. We argue that our account has a distinctly epistemic character, irreducible to the accounts of collective moral responsibility. As such, it fills an important gap in the literature on epistemic responsibility. In contrast to previous accounts of epistemic responsibility, which have been concerned with the conditions of responsible belief formation and holding, our approach concerns responsibility for other kinds of performances, specifically those aimed at preventing epistemic harms.

1 Introduction

Consider the following two scenarios.

Abandoned-research. Imagine scientists working in a certain branch of medicine, trying to find a cure for a certain disease. Since there are two main hypotheses for what could be causing the disease, some scientists are pursuing the line of inquiry based on hypothesis H_1 while others follow hypothesis H_2 . Due to a random course of events (such as a number of scientists having retired from academic research, or having moved on to a different research project due to personal career goals) H_1 is dropped. A few decades later it turns out that the

abandoned hypothesis in fact offered the right track to finding the cure for the disease, and the scientific community ended up side-tracked on a sub-optimal line of inquiry for a long period of time.¹ In such a case, who (if anyone) is to blame? Moreover, do contemporary scientists have any duties aimed at lowering the risk of a premature abandonment of fruitful hypotheses?

Biased. There has been a growing body of evidence that some practices employed in a certain field of scientific research frequently fall prey to various types of biases, lowering the overall reliability of published results. If I am a scientist working in this field and aware of the problem, yet I know that I can't change the entrenched practices on my own, do I have any moral and epistemic duties to address this problem? And conversely: if I fail to act towards resolving the issue, am I morally and/or epistemically blameworthy for what is happening?

Each of the above cases concerns a situation in which the given scientific community is facing a threat of epistemically harmful effects (or epistemic harms for short). Moreover, each of the epistemic harms is such that it could be prevented only by a group of scientists rather than by any individual member of the given scientific community. While various scientific institutions act as organized groups that aim at both forwarding epistemic goals and preventing epistemic harms, the above cases illustrate problems that typically don't fall under the jurisdiction of an existing institution, but require a joint action of scientists working in the given domain, where they are unorganized with respect to the given issue. Different theories of collective moral responsibility have been developed to account for moral duties of unorganized groups in such circumstances.² The basic intuition they aim to address is that individuals should join forces in order to prevent the given harm. But what if the given harm is such that it has distinct epistemic features (e.g. in virtue of impeding our knowledge acquisition)? What if, for instance, the scientists in *Abandoned-research* are cosmologists, so that abandoning a fruitful hypothesis may have no major social or ethical implications, though it may have epistemic ones? Similarly we may ask: is the nature of the duties we associate with *Biased* primarily moral, or also (and equally) epistemic?

In this paper we propose an account of collective epistemic responsibility which provides the normative basis for answering these questions. Our approach is inspired by Hindriks' (2018) account of collective moral responsibility. It is rooted in the idea that when a harm can only be prevented by group action, this creates a specific duty for each individual in the group: namely, a duty to join action in order to prevent the harm.

¹This story is pretty much what happened during the twentieth century research on peptic ulcer disease, see Šešelja and Straßer, 2014b.

²We will provide an overview of these accounts in Section 2.

Our account offers three significant novel contributions. First, it distinguishes between unorganized groups and what we call mobilized groups. This helps motivate and justify two-part views like ours and Hindriks'. Second, we will argue that collective *epistemic* responsibility has a distinct normative character. Thus, our account complements accounts of collective *moral* responsibility. Third, our approach fills an important gap in the literature on epistemic responsibility since, unlike previously proposed accounts, it provides a normative framework for evaluating cases of epistemic harms that can be prevented only by a collective action. This last point requires some clarification.

The literature on epistemic responsibility traditionally concerned appropriate belief-formation (e.g., Code, 1987; Hieronymi, 2008; Kornblith, 1983; Miller and Record, 2013; Robitzsch, 2019). In particular, it concerned whether subjects acted appropriately in response to what evidence (or other epistemic resources, like methods, processes, or virtues) they had available to them when forming beliefs. Subsequently, accounts of collective epistemic responsibility have generally concerned collective belief-formation and knowledge acquisition (e.g., Corlett, 2001, Rolin, 2008; Rolin, 2016, Millar, 2020). The theory we propose here is distinct from these traditional projects in two important ways. First, our project concerns not belief-formation, but other kinds of epistemic performances and activities. In particular, we focus on actions at earlier stages of inquiry than belief-formation. Second, our focus is on norms that prevent a failure to meet epistemic goals, which are cases of epistemic harm. As such, our account is *preventionist* in character. The main reason why traditional accounts of collective epistemic responsibility need to be supplemented by a preventionist one is that preventing a failure to achieve epistemic goals may require different norms than those typically guiding belief-formation. In other words, a group of agents who have engaged in responsible belief-formation (or knowledge acquisition) may still fail to engage in the prevention of pending epistemic harms.

Our example *Abandoned-research* is a case in point: while scientists pursuing H_2 may hold an adequate epistemic attitude towards their hypothesis (e.g. treating it as only one among many fruitful hypotheses), and even consider H_1 worthy of pursuit, traditional accounts of epistemic responsibility do not address the question: what kind of responsibility do these scientists have with respect to the threat of an epistemic harm posed by the abandoning of H_1 ? Similarly, if we examine the case retrospectively, traditional accounts don't tell us who (if anyone) is to blame for the given abandonment and why.

Here is how we will proceed. In Section 2, we offer some background on the notion of collective moral responsibility, and discuss the duty to join forces. In Section 3 we introduce our account of collective epistemic responsibility and show how it addresses the cases we began with. In Section 4 we explain how our account is genuinely epistemic in character. In Section 5 we conclude the paper by showing how our account fruitfully applies to some additional cases.

2 Collective Moral Responsibility

Often, people talk as though groups have obligations, or are responsible for things. We hold governments and corporations accountable for their actions. Exxon-Mobil is responsible for various oil spills. BP is obligated to clean up the Gulf of Mexico. This suggests that there are group obligations and responsibilities. A common way to classify groups that bear such obligations is as follows. The first are organized groups with explicitly specified structures and decision procedures, like corporations and governments. The second are persistent but unorganized social groups, e.g., races, genders, and nationalities. Finally, and most controversially, there are random collectives (Held, 1970): groups of people who are only connected by the relevance of some problem or task, e.g., passengers in train car 6745, or beachgoers in Asbury Park on July 23rd.

There is something puzzling about attributions of collective responsibility. Generally, we only hold moral agents responsible for their actions. So, some philosophers have argued that organized groups like Exxon-Mobil and the United States count as group agents. But this response is not particularly plausible as an explanation for why persistent social groups bear responsibility, and it is even less plausible as an explanation for attributions of responsibility to random collectives. Yet there are a variety of cases where it is intuitively plausible that even fleeting, unorganized groups do bear collective responsibility. Many such cases involve collective action problems, what Hindriks calls “collective harms” (Hindriks, 2018). In this section, we will discuss desiderata for a theory explaining the intuitive appeal of assigning collective responsibility in such cases.

There are two important distinctions to make when talking about responsibility. First, there is a distinction between accountability and what we might call positive responsibility (Williams, 2008). Sometimes, when we say an agent is responsible, we mean that it is appropriate to hold them accountable for something: it is appropriate to praise or blame them. The antonym of this kind of responsibility is “not responsible.” On the other hand, sometimes we say someone is responsible because they have met their obligations, and are praiseworthy. The antonym for this positive responsibility is “irresponsible.” Epistemologists, especially responsibilists, have been interested in both kinds of responsibility (Baehr, 2011; Williams, 2008; Zagzebski, 1996). In this paper, we will be concerned with responsibility as accountability. The second relevant distinction concerns backward-looking responsibility, as associated with praise and blame, and forward-looking responsibility, which is associated with obligation and remediation (Smiley, 2017). However, our view, like many views in the contemporary literature, will seek to apply coherently to both kinds of responsibility. So, we will largely set the distinction aside in what follows.

2.1 Desiderata for a Theory of Collective Responsibility

Consider the following case (adapted from Björnsson, forthcoming):

Beach: A group of twelve children are swimming in the ocean. Three of the children have brought a parent with them. Suddenly, the wind changes and begins sweeping the children out to sea. Each adult only has time to save one child by swimming. However, there is a boat nearby that can be operated by two adults. With the boat, all the children can be saved.³

Intuitively, the parents have an obligation to save all the children. The only way to save them, however, requires a coordinated collective action. No individual adult can save all the children. Only the coordinated action of two adults together can operate the boat. Hence, the children being swept out to sea is a collective harm: it can only be prevented by collective action. Moreover, the parents do not form any kind of organized group; they may not even know one another. The only connection between them is that they happen to be at the beach at the relevant time. They are thus an unorganized group, what Held calls a “random collective.” (1970, p. 475). Despite each individual parent’s inability to save the children, and despite their lack of organization as a group, it is still intuitive that parents are obligated to save all the children. This is an instance of what we can call the *primary intuition* about collective responsibility for unorganized groups.⁴ Of course, the structural similarity of this case with the two cases we began with is no coincidence.

A theory of collective moral responsibility must explain the primary intuition.⁵ However, this isn’t the only requirement for such a theory. Schwenkenbecher suggests a variety of desiderata for an account of collective responsibility (2018, pp. 111-112). For one thing, a theory needs to explain additional intuitions, e.g., that each individual has responsibilities in such cases, and that their (other) individual responsibilities can sometimes come into conflict with the group responsibilities. In *Beach*, the individual parents each have an obligation to save their own children, and this might conceivably conflict with the group’s duty to save all the children. At the same time, each parent also seems to have an individual duty to contribute to the collective action solution. Call these, and related intuitions, the secondary intuitions.

In addition to explaining intuitions, a theory of collective responsibility should also cohere with accepted moral principles. Or at least, it speaks in favor of such a theory if it does not propose ad hoc exceptions to these principles. Hindriks helpfully identifies four relevant conditions on responsibility which are commonly accepted principles of this sort (2018, p. 206):⁶

³Virginia Held contributed a variety of cases like this to the literature in her (1970) paper. Her discussion there inspired much of the subsequent literature on collective responsibility for unorganized groups. For another influential discussion, see Parfit (1984).

⁴We are following Bjornsson (forthcoming) and Schwenkenbecher (2018) here in emphasizing this.

⁵Some philosophers are skeptical of appeal to intuition. However, it is commonly accepted in the ethics and epistemology literature that intuitions provide at least defeasible evidence for a theory, and that they serve as part of the data to be explained by such theories. We will follow that convention here.

⁶See Hindriks, 2018, p. 206 for a discussion on possible variation on each of the listed

1. *The agency condition*: Only full moral agents can bear responsibility. A full moral agent is normatively competent in the sense of being receptive and responsive to epistemic reasons (thus conforming to Fischer and Ravizza's (1998) reason-responsive account of responsibility).
2. *The causal condition*: the agent is able to prevent the harm. In the collective case, this might be in virtue of a joint action with other agents.⁷
3. *The epistemic condition*: the agent has a justified belief about the existence of the pending harm and the likelihood of the success of preventing it (potentially by means of a collective effort), or she is in an epistemic position such that she is able to have a justified belief about this.⁸
4. *The no-defeaters condition*: The agent does not have defeating evidence that provides an excuse or a justification for not fulfilling the duty (for example, evidence showing that preventing the given harm by a collective effort is in conflict with another epistemic or moral duty).

As will become clear, explaining the primary intuition without violating these principles is difficult. In particular, it is hard to explain the intuition in cases like *Beach*, as there are no agents who satisfy both the agency and causal conditions. The group is not an agent, while no individual parent can save all the children.

In addition to explaining the intuitions and cohering with commonly accepted moral principles, there are several other desiderata. Many philosophers have thought that a theory of collective responsibility for unorganized groups should be action-guiding. Moreover, accepting the theory and following its dictates should lead to moral improvement. The thought is that a theory which vindicates the moral responsibility of groups will enable us to argue that people are required to take part in collective action solutions such as preventing climate change. To this end, a theory should not only posit collective responsibility of groups, it should explain how individuals' responsibilities are to be derived from (or at least related to) collective responsibility.

Finally, theories of collective responsibility should be evaluated in part based on consideration of more generic explanatory virtues. A theory will thus be better insofar as it is: more parsimonious, because it posits fewer new entities or novel principles; has greater explanatory power, because it explains more of the data in question; or it is more unified (i.e., consilient), because it provides an explanation for evidence from diverse sources, or coheres well with other explanations which do so. These desiderata were proposed for theories of moral responsibility. However, we suggest that the same desiderata, suitably modified, are relevant for evaluating theories of collective epistemic responsibility.

conditions. Schwenkenbecher identifies the same principles, though describes them somewhat differently (2018).

⁷Schwenkenbecher calls this the *capacity principle* (2018, p. 111), while others have given similar conditions in terms of *ability* (Pinkert, 2014; Wringer, 2019).

⁸Hindriks formulates this condition in terms of an agent knowing of the pending harm, but adds that this may be a too strong requirement. The second disjunct points to the importance of agents being epistemically responsible with respect to the formation of such beliefs.

2.2 The Duty to Join Forces

There are two important choice points for a theory of collective responsibility. The first is reducibility. Reductive accounts explain away intuitions about collective duties by reducing them to duties of individuals. Non-reductive accounts suggest that collective responsibility is irreducibly ascribed to a collective as such. The second choice point concerns how to deal with the agency condition.⁹ Conservative theories attempt to accommodate the agency condition. They seek to explain away the primary intuition by appeal to responsibilities of either individual or group agents.¹⁰ Revisionist theorists, following Held (1970), argue that being a full-fledged, reasons-responsive moral agent is not a requirement for bearing responsibility. Wringer (2016; 2019), for instance, takes the primary intuition to be strong motivation, by itself, for discarding the agency condition. Finally, joint (or shared) theories ascribe moral obligations and responsibilities only to individual agents. However, they suggest that the *content* of those responsibilities is distinctively shared, as they call for irreducibly joint actions.¹¹

Each of these choices introduces problems for meeting the desiderata. Reduction requires denying the primary intuition, while anti-reduction requires explaining who or what has the non-reducible responsibility. Conservatives often deny the primary intuition. Revisionists must deny the very plausible agency condition. And joint/shared theorists must motivate a variety of novel explanatory machinery to vindicate the idea that responsibilities can be both aimed at individuals and still genuinely shared.

Hindriks (2018) proposes the *Duty to Join Forces* theory as a hybrid of the above approaches, which avoids their pitfalls while keeping their strengths. The central idea of his account is that collective responsibility is explained as a duty in two stages: first, as a duty to join forces, and second, a duty to prevent the harm. Hindriks' point of departure "is the claim that a random collective has a duty to prevent an outcome only if enough of its members are ready to suitably combine their preventive efforts. Furthermore, such a collective often acquires this duty only after a sufficient number of members have been mobilized." (p. 205).

Thus, the view proposes two stages of responsibility that together comprise the duty to join forces. The stages are:

1. Mobilize others
2. Collectively prevent the harm.

⁹Here, we are following the taxonomy of (Schwenkenbecher, 2019).

¹⁰Feinberg offers a conservative view that denies the primary intuition (1968). Other traditionalists are happy to talk about obligations of groups if they show appropriate cohesion. This cohesion might be a matter of group organization which allows for group agency (List and Pettit, 2011; Pettit and Schweikard, 2006; Tollefsen, 2015), or of solidarity or identification (Darby and Branscombe, 2014; McGary, 1986).

¹¹Joint views are proposed by, Björnsson (2014), Miller (2015b), Pinkert (2014), and Schwenkenbecher (2019), among others. Schwenkenbecher, for example, argues that joint responsibilities exist when agents are required to "we-reason." This requires each individual to reason about what "we" should do, as a group.

The first stage of mobilizing others consists in a responsibility of individuals. This is a responsibility of each member of the group to communicate with the other group members and convince them to join the effort necessary for preventing the harm. Hindriks calls this a *lateral* duty, as it requires not that the agent prevent the harm (which is the ultimate goal), but that the agent communicate and convince other members of the group. In addition, individual group members have a responsibility to be receptive to others who attempt to mobilize them. The first stage is successful if an adequate number of group members have been suitably mobilized to join the collective effort.

The second stage is a conditional norm: it only occurs if the first is successful. If this condition is satisfied, then, the collective as a whole has an obligation to prevent the harm. The collective now has a duty to engage in the joint action needed to avert the bad consequences.¹²

We can illustrate the duty to join forces by appeal to *Beach*. First, the parents have a responsibility to mobilize: to come to agreement about what they should do. In this case, each should communicate with the others about the presence of the boat, and how they can work together to use the boat for rescuing all the children. Once there is agreement among enough of them about what course of action to take, they have successfully mobilized. This activates the second stage: the mobilized group of parents now has a responsibility to use the boat to save all the children. At the time they have an obligation (forward-looking responsibility). If they fail to do so (through inaction, say) they can be held accountable (backward responsible) for this failure. If both stages are completed successfully, then the parents will have fulfilled their duty to join forces, i.e., their collective duty. Thus, the account explains the primary intuition as applied to this case.

The status of a random collective after enough people have been successfully mobilized to engage in collective action (other things being equal) is left somewhat unclear by Hindriks' discussion. Hindriks highlights the difference in ontological status between group agents and groups that have joined forces. However, we think there is also an important distinction between the random collective before and after joining forces (Hindriks, 2018, p. 211). We will return to this below.

Hindriks' account fares well regarding the desiderata introduced above (section 2.1). It explains the primary intuition. It also explains the secondary intuitions regarding individual duties: individuals have duties to mobilize others, and then to engage in the joint action required to prevent the harm. It coheres well with the previously accepted moral principles, with the notable exception of the agency condition, as we will discuss below (section 3.3). It is parsimonious in that the only novel things it proposes are the collective obligations which are needed to satisfy the primary intuition. It does not require appeal to novel types of reasoning, or commitment to anything particular about

¹²Collins (2013) proposes an account of collective responsibility in terms of a duty to incorporate into a group agent. However, she admits that in cases like *Beach* it is implausible to think that the collective needs to form an organized group like a corporation in order to save the children.

moral motivation (as do Schwenkbecher and Björnsson, respectively). It does not add anything ontologically substantial to suggest that individuals have obligations to mobilize others, as individual responsibilities are commonly accepted, and in any case are necessary to vindicate the secondary intuition.

3 Collective Epistemic Responsibility

In this section, we offer a theory of collective *epistemic* responsibility. Our theory is a hybrid account inspired by Hindriks' Duty to Join Forces.

3.1 The epistemic duty to join forces

We start from the idea that there are certain epistemic harms that can be prevented only by a collective of individuals, rather than by any of its members on their own. We will call such effects *collective epistemic harms*. Moreover, such harms may occur in situations where there is no existing organized group that could prevent it. As the examples at the beginning of this paper illustrate, situations in which scientists working in a given domain may end up acquiring knowledge inefficiently and accepting false hypotheses, where it would take more than one individual to prevent such a harm, are not hard to imagine. Let's specify these notions:

Epistemic harm: a harm affecting the epistemic status of a subject, group of subjects, or epistemically important social system.¹³

This characterization of epistemic harm is quite generic: it intentionally leaves open the substance of an epistemic harm. This allows our account to be neutral between a wide variety of different epistemological theories of epistemic value, justification, scientific progress, etc. Our theory of collective epistemic responsibility thereby applies to a wide range of epistemic views. For example, an easy way to make the notion of epistemic harm concrete and substantive is to adopt a veritistic, reliabilist account (Goldman, 1979a, 1986). On such a view, a harm involves causing subjects to hold false beliefs, or undermining the reliability of their individual or social methods of belief formation.¹⁴ Alternatively, in providing a theory of epistemic injustice, Miranda Fricker offers a pluralistic view which includes a variety of ways to be epistemically harmed (Fricker, 2007).¹⁵

¹³For more on systems epistemology, see Goldman, 1999.

¹⁴There are a variety of competing views about the nature of epistemic value and the goals of inquiry, any of which will be compatible with our account. The view one chooses will help determine what counts as a harm due to impeding such goals. Besides truth, the goals might be knowledge (Williamson, 2002), understanding (Elgin, 2017; Kvanvig, 2003), problem solving (Laudan, 1977), etc. In any case, the goal of any inquiry likely needs to be constrained by the question guiding that inquiry (Friedman, 2017; Hookway, 2008; Khalifa, forthcoming; Millson and Khalifa, 2020), which will also constrain what can count as a harm.

¹⁵An epistemic injustice involves epistemic harms which are unjust, generally because of oppressive social structures. For more on this and related ideas, see (Dotson, 2011; Fricker, 2007; Kidd, Medina, and Pohlhaus, 2017; McKinnon, 2016; Medina, 2013).

We are interested in a specific type of epistemic harm: the kind that can only be prevented by group action. Thus:

Collective epistemic harm: an epistemic harm that can be prevented only by a joint effort of several individual agents, rather than by any single individual on their own.

The existence of such harms, as illustrated by our motivating cases from section 1, shows the need for a theory of collective epistemic responsibility that applies even to unorganized groups. Each of those cases elicits an intuitive judgment analogous to the primary intuition in *Beach*. For instance, in *Abandoned-research*, there is a collective epistemic harm that results from abandoning a live hypothesis which turns out to be true. The intuition is that the collective, comprised of the scientists working in the relevant branch of medicine, is responsible for negligently abandoning an important (and in this case true) theory. Call this the *primary epistemic intuition*. There is also a secondary intuition: that individual medical researchers who are members of the collective bear their own, individual responsibilities for failing to contribute to a joint solution to the collective epistemic harm. Moreover, there are other secondary intuitions: for example, that individual scientists also bear individual responsibilities towards their own respective inquiries, which may be incompatible with the abandoned hypothesis, giving rise to possibly conflicting obligations.

The existence of collective epistemic harms, and the primary and secondary epistemic intuitions, leaves social epistemology with a problem similar to that of collective moral responsibility in ethics. This is one upshot of our discussion worth highlighting, even for those who will disagree with our specific account: there is a problem that requires a solution. The desiderata for solving this problem, we suggest, are analogous to those discussed above in section 2.1.

In view of this we propose a theory of collective epistemic responsibility based on a two stage duty:

Epistemic duty to join forces: an obligation of an unorganized group to prevent a collective epistemic harm. It consists of the following sub-duties:

- D1 a duty of individuals to approach other relevant agents raising awareness about the epistemic harm, expressing willingness to prevent it, and encouraging others to do the same;
- D2 a duty of those who have fulfilled D1, and thereby formed a mobilized group, to prevent the epistemic harm.

The epistemic duty to join forces (EDJF) is a two stage-view, just as in Hindriks' account of collective moral responsibility. It is a conditional norm: D2 is triggered when for a sufficient number of involved individuals D1 is successfully fulfilled. The first stage is a responsibility of individual agents to mobilize others. This requires communication and organization. If D1 is successful, then the mobilized collective is responsible for preventing the harmful outcome, first in *prospective* sense, meaning there is a collective epistemic duty to follow

the norm, and subsequently in the *retrospective* sense, meaning the group is epistemically praise- or blame-worthy.

Note that if there is an unorganized group in a position to prevent a collective epistemic harm, but the members of the group fail to fulfill D1, then D2 is not triggered. In this case, individual members of the group will be blameworthy for failing to join forces (for failing D1). But since the collective has never been adequately mobilized, the collective never obtained a duty to prevent the harm. Thus, according to EDJF, duties to prevent collective harms remain conditional: the collective duty only obtains once the individual duties to join forces have been fulfilled. This has the benefit of ascribing duties in such cases only to moral agents who fulfill the conditions of all four standard principles of responsibility (described in section 2.1). Only at the second stage is their responsibility born by the collective. If the collective is never suitably mobilized to prevent the harm, i.e., there is no joining of forces (we will shortly give an example of what that means in the context of our cases), then D2 simply never obtains. Nonetheless, EDJF does predict that members of the collective bear epistemic responsibility in such a case, and will thereby be epistemically blameworthy as individuals for failing to fulfill D1. Moreover, the view suggests a sense in which the unorganized collective is responsible: all of its members have a duty to join forces. This helps to vindicate the intuition that there is blameworthy action occurring in such a case. D1 is thus a *reducible* group responsibility. The sense in which the group is responsible is reducible to the fact that its members are.¹⁶

A group which has successfully joined forces, and met its D1 duty, must be importantly different in its structure than it was prior to mobilization. This new structure will generally fall short of what is required for group agency. Still, the reason why it seems plausible to attribute a duty to the group once it has been mobilized is that it has some new kind of status. Otherwise, it would not help to require that the first stage be completed before the truly collective duty can be ascribed to the group.

We propose that groups which have completed their stage one duty to join forces be called *mobilized groups*. Such groups have more organizational structure than a truly unorganized group. Moreover, they may have different members than the initial group. For instance, if two of the three parents in *Beach* succeed in joining forces (or two-three scientists in *Abandoned-research*), they are now a mobilized group. They now have a second stage duty to save the children (i.e. to prevent the abandonment of the given hypothesis). This mobilized group has fewer members than the unorganized group.

Mobilized groups lack organizational structure or long-term persistence. Still, it is plausible that they have adequate cohesion to warrant ascription of collective responsibility. Compare these to ascriptions of moral responsibility to persistent social groups. For instance, many have found it plausible that white

¹⁶Of course, this is compatible with some of the members having excuses. What is required is that some significant subset of the group has such duties. What subset this is will depend on what is required in order to avoid the epistemic harm, but it will need to be more than a single individual.

Americans owe reparations for slavery and racial oppression.¹⁷ Regardless of whether collective responsibility obtains in this specific case, its plausibility suggests that persistent social groups are potentially responsible for things, despite lacking the kind of organizational structure required for group agency. Similarly, we suggest that mobilized groups are potentially morally responsible, despite lacking full-blown agency. Recognition of mobilized groups, and their status as having a moderate level of organization, is a novel aspect of our account.

Recognition of mobilized groups allows EDJF to explain the primary intuition, while meeting all four of the conditions on responsibility discussed above (Section 2.1). We will return to this point below, after discussing how EDJF handles our motivating cases.

3.2 Back to the examples

EDJF can be illustrated by appeal to the examples from section 1. We first consider *Abandoned-research* as well as several variants of this example. This will illustrate the nuanced predictions and evaluations that the epistemic duty to join forces allows. It will also show how the account handles both the prospective and retrospective kinds of responsibility.

In *Abandoned-research* a plausible and potentially true theory, H_1 , is dropped from active research. The question is who, if anyone, can be considered blameworthy for the premature abandonment of the fruitful hypothesis? The original version of *Abandoned-research* concerns retrospective responsibility: H_1 has been abandoned already, and the question is who is culpable. EDJF analyzes this case by appeal to its two stages. In particular, D1 is active as a requirement for the group. The research community has a duty to join forces in order to prevent the epistemic harm from the lack of research on H_1 . The individual members of the community had a duty to communicate with one another about avoiding the harm ahead of time. This could have been accomplished by an agreement regarding this particular theory, in the form of individuals volunteering to rescue H_1 and pursue it. Or it could take the form of a more general solution, e.g., setting up a special research fund to ensure that such theories are generally rescued. Either option would require communication and agreement on some reformulation of the organization of the medical research community. Since researchers in the community failed to engage in this communication, or make any attempt to get others to join forces, the community has failed to fulfill D1. It is retrospectively responsible for this failure.

Consider a variant of the case:

Abandoning: This case is precisely like *Abandoned-research*, except that instead of ten years later, we consider the situation just before H_1 is abandoned. It is apparent that this abandonment will occur

¹⁷For an overview of literature on this topic, see (Boxill, 2016). For a discussion particularly relevant to our claims about cohesion, see (McGary, 1986).

soon, as H_1 's main proponents are retiring and no one intends to step up to take their place.

In *Abandoning*, EDJF suggests that there is a prospective responsibility of type D1: the research community is obligated to communicate about how to save H_1 . Each member is required to approach others and attempt to convince them that H_1 should be rescued. Moreover, each is responsible for being receptive to such entreaties. Epistemic responsibility requires trying to recruit others, and being willing to be recruited. If they fail in this duty, the situation will end up precisely as it does in *Abandoned-research*. The community will be (retrospectively) responsible for its failure to join forces, in virtue of its members' failure. However, if the community in *Abandoning* fulfills its obligation to join forces, it will thereby obtain a D2 obligation: to prevent the abandonment of H_1 . This suggests two more variants to the case:

Rescued: This case is just like *Abandoning*, but the community has in fact succeeded at joining forces. There is widespread communication and agreement that H_1 must be saved. The community agrees that certain members will be assigned to pursue it, they do, and the community arrives at the truth much more quickly.

Failed rescue: This case also proceeds just like *Abandoning*, and here the community also succeeds at joining forces, achieving widespread agreement that H_1 must be saved. However, after this agreement is reached, the community spends too long deliberating about which researchers, in particular, should pursue the theory. Research is stalled, and H_1 is effectively abandoned anyway.

In *Rescued*, EDJF is completely fulfilled. The community successfully fulfills its D1 obligation to join forces and mobilizes, thereby obtaining a D2 obligation to prevent the collective epistemic harm of abandoning H_1 . Finally, D2 obligation is fulfilled when members of the community agree to take on responsibility to pursue the theory. This illustrates successful fulfillment of prospective responsibility. Moreover, this version of the community subsequently is responsible for preventing the harm, and is thereby praiseworthy.¹⁸

Failed rescue, however, illustrates a community which fulfills D1, but still fails to meet its collective obligation. In this case, the EDJF suggests that the community has failed a collective duty. It is genuinely collectively responsible for failing to prevent the harm (the abandonment of H_1). It is culpable, qua group, for this failure.¹⁹

¹⁸Note that according to our account the conditions for triggering EDJF are independent of whether the given hypothesis is actually true or not. Even if H_1 was rescued but turned out to be false after all, triggering EDJF was warranted as long as H_1 appeared worthy of pursuit so that abandoning it indicated a threat of an epistemic harm (for discussions on such situations see Fleisher, 2018; Šešelja, Kosolovsky, and Straßer, 2012; Šešelja and Straßer, 2014a).

¹⁹Of course, a group might have an excuse a failure to fulfill D2, in which case they will not be culpable for the failure. *Failed rescue* makes it explicit that it was inappropriate deliberation that leads to the failure, which suggests no such excuse, and furthermore we can assume there is no other excuse available.

These variant cases show the flexibility of EDJF in applying to a variety of situations. Moreover, they show that the theory provides us with a way of diagnosing normatively relevant differences between circumstances. The group in *Failed rescue* does seem to have done better than the one in the original version of *Abandoned-research*: it has at least recognized the problem and joined collective forces with the aim of solving it. It just fails this final step. At the same time, the *Failed rescue* group has a duty that the original group lacks. This makes sense, since they have obtained an ability the original group lacks: the ability to prevent the harm. A disorganized collective is not able to solve such a problem, while a mobilized collective is. This allows the view to prescribe duties only to entities that have the capability of fulfilling them, as required by *the causal condition*. This last point helpfully motivates the discussion in the next section. There, we will consider how EDJF, and its handling of the cases at issue, fares with respect to the desiderata raised above in section 2.1.

EDJF's handling of our second case, *Biased*, functions in the same way, so a brief sketch should suffice. Recall that this case concerns a random collective comprised of scientists in a research field with biased methodology. Continuing with this biased methodology represents a collective epistemic harm: it is harmful because it impedes scientific discovery, and can only be prevented by collective action. The question here is whether such a group bears a collective epistemic responsibility to change this standard. EDJF suggests that this group, too, has an epistemic duty to join forces. Prospectively, it has a D1 reducible duty, as each scientist will have a duty to communicate about the problem and be receptive to joining a collective effort to solve it. If enough of the community agrees to change, the newly mobilized collective will thereby incur a D2 obligation to change the methodological practices of the field. If they fail their D2 obligation, the collective will be retrospectively responsible for this failure (unless they have some excuse). Moreover, if the group fails to join forces, fails to meet their D1 obligations, then the group will be retrospectively responsible for failing to join forces, in virtue of the responsibility of the individual scientists for failing to do so.

3.3 Desiderata Revisited

How does EDJF, and its proposed analysis of the cases, stack up regarding the desiderata for an account of collective responsibility? In this section, we evaluate the theory in light of these goals, before moving on to argue in Section 4 that the responsibility we have identified deserves to be called “epistemic.”

Primary and secondary intuitions. First, let's consider the primary intuition as it applies to our instigating cases. In *Abandoned-research*, the primary intuition is that the random collective comprised of researchers in a medical field had a collective responsibility to prevent the abandonment of hypothesis H_1 . The variant cases considered in the last section showed that there is both a prospective and retrospective version of this intuition. EDJF explains this intuition in each case. The collective has an epistemic duty to join forces. If

they fail to do so, they retrospectively bear responsibility for failing to fulfill this duty. This failure can come at either stage. The community would also bear responsibility (and be praiseworthy) if they had tried and succeeded in joining forces to prevent the abandonment. EDJF has a disjunctive strategy for explaining the intuition. Either the responsibility of the group is reducible to the individuals' responsibility to join forces (D1); or the responsibility is irreducible and applies directly to the newly organized collective (D2). According to either option, EDJF licenses the judgment of collective responsibility that constitutes the primary intuition.

Next, we consider secondary intuitions. We noted above two important secondary intuitions for a theory to address. First, that individuals also have responsibilities to contribute to preventing collective harms. Second, that other individual responsibilities can come into conflict with the group responsibilities. The epistemic duty to join forces offers explanations for both of these secondary intuitions. According to EDJF, individuals have D1 responsibilities to communicate with others, attempt to enlist them in a coordinated effort, and to be receptive to others attempts to do the same. These are individual responsibilities to contribute to the collective action solution. At the same time, individuals may have other obligations. For instance, in *Abandoned-research*, a researcher may have a responsibility to continue work on H_2 in light of the funding they have received, or because their lab is best positioned to test that hypothesis. This is a conflicting responsibility, because it at least suggests that the researcher should continue such work at the expense of spending time enlisting others to save H_2 . This is structurally similar to the conflicting moral responsibility in *Beach*, where the parents responsibility to their own children conflicts with the responsibility to join forces to save the group. Which obligation is the more important for an individual to follow will depend on the details of the case. But any plausible theory will need to allow for such a conflict, and EDJF does so.

Coherence with the principles of assigning responsibility. We now turn to the harmony of EDJF with existing principles of responsibility. The primary ones at issue are the four conditions (on both retrospective and prospective responsibility) from section 2.1: the agency, causal, epistemic, and no-defeaters conditions. We have previously seen their motivation in the context of moral accounts of responsibility. Here we will start by assuming that prima facie they also hold in the context of epistemic harms and epistemic responsibility. Since the final three of the above conditions are reasonably straightforward for EDJF, we treat these first. The agency condition is more complicated, and we consider it last.

The no-defeaters condition requires no modification or special consideration as applied to collective epistemic responsibility. EDJF is perfectly compatible with a recognition that an agent (or group of agents) may have other reasons, or excuses, for not preventing a harm. Both the D1 and D2 obligations are compatible with this recognition, and are potentially defeasible by other epistemic

or moral considerations.

The causal condition (also called the ability or capacity condition) requires that an agent bears a responsibility to ϕ only if they are able to ϕ . That is, the agent must have causal efficacy with respect to what they are required to do.²⁰ In EDJF's analysis of *Abandoned-research*, the causal condition was fulfilled for D1, since each individual scientist is plausibly able to communicate about joining forces (by email, in conferences and publications, etc.), and able to signal willingness to join forces when others initiate such communication. The condition is also fulfilled for D2, since upon joining forces the scientists would have been able to prevent the abandonment of H_1 by agreeing to an adequate division of labor in the community.²¹ Thus, EDJF fits well with this accepted principle, requiring that the causal condition be met in cases of responsibility.

The next one is the epistemic condition. Recall that this condition requires that to be responsible for preventing some harm, an agent must know (or be in a position to know) about the existence of the pending harm (or the risk of such a harm occurring).²² Moreover, they must know of some reasonably probable way of preventing it. Some views have suggested justified belief in place of knowledge, but either way, justified belief would seem to be required.

The epistemic condition may initially seem worrisome: one might wonder whether there is potential circularity in applying an epistemic condition to an account of epistemic responsibility. In other words, the worry is that epistemic responsibility is required for justified belief (and knowledge), while justified belief is required for epistemic responsibility. However, this worry is easily dispelled by recognizing the difference between the notion of epistemic responsibility associated with responsible belief formation and belief holding (Code, 1987; Hieronymi, 2008; Kornblith, 1983; Montmarquet, 1993; Robitzsch, 2019; Rolin, 2008; Williams, 2001; Zagzebski, 1996) and the notion associated with the prevention of epistemic harms, which we are concerned with. If the epistemic condition requires a responsibly formed belief (a controversial claim), this means a positively responsible belief: one that is formed in a way appropriately responsive to evidence, and where evidence is gathered properly, etc. This is a separate kind of responsibility than described by EDJF, which is an account of preventionist responsibility. However, even if it turned out that the being positively responsible requires being responsible in a preventionist way, there is no special problem for our account. Being responsible (or justified) in believing one proposition often involves a prior requirement of responsibility (or justification) in another proposition. Any problems this causes (e.g., the traditional regress of justification problem) are in no way specific to our account.

With the above ground-clearing done, it should be easy to see how EDJF

²⁰This causal efficacy need not require that the agent be able to ϕ at will. The ability to ϕ reliably enough when one tries is adequate in most cases.

²¹Implementation of this plan would include, for instance, mentioning all the promising hypotheses in the articles and textbooks, encouraging others to embark onto less pursued paths, etc.

²²For instance, in *Abandoned-research* it is enough if the scientists know there is a substantial risk of an epistemic harm happening, rather than a full certainty that it will happen.

can cohere with the established principle. According to EDJF, D1 obligations require individual agents to be in a position to know about the pending collective epistemic harm, and potential steps to a solution. In *Abandoned-research* and its variants, this means being positioned to know about the pending harm of abandoning H_1 , and having some idea of how to begin to solve the problem.²³ D2 obligations also require that the group be in a position to know about the harm. But for D2 responsibilities to be operative, the D1 responsibility to communicate must already have been fulfilled. This makes it plausible that a group that has joined forces will have a group belief about this in a very straightforward sense: all the members who have joined together will know about the harm and potential solution since they will have talked about it. Thus, the epistemic condition is a reasonable requirement for both stages of EDJF, and the theory is thereby compatible with the established principle.

The situation with the agency condition is more complicated. The agency condition requires that only full moral agents—or for our purposes, *intentional epistemic agents*²⁴—those that are normatively competent, can be the bearers of (epistemic) responsibilities and obligations. This principle has traditionally been the sticking point for theories of collective moral responsibility. It is plausible that organized groups, like corporations, have the requisite decision-making structure and sensitivity to reasons required for counting as full moral agents. However, *ex hypothesi* unorganized groups do not constitute group agents.

Advocates of the idea that random collectives can be morally responsible for preventing collective harms have argued that the agency condition is false, at least as it is traditionally understood. Wringe (2019), for instance, takes the primary intuition to be powerful evidence that unorganized groups can bear responsibility. He argues that this by itself shows the falsity of the agency condition, and suggests instead that any entity capable of preventing a harm meets whatever conditions are necessary for responsibility. Hindriks (2018) argues that the agency condition is too restrictive. It was intended to rule out holding animals, small children, and other non-competent agents as responsible. It was never meant to apply to the collective cases. He opts for a revised version which allows that either full moral agents or groups made up of such can bear responsibility. Analogously, we could allow that either intentional epistemic agents or groups made up of such bear epistemic responsibility. If either Wringe or Hindriks' arguments are compelling, EDJF is thereby made harmonious with the resulting set of principles.

It is worth noting, however, that the original, restricted agency condition is even less plausible in the epistemic case than the moral one. The notion of collective epistemic agency is no novelty in the philosophical literature. For example, according to Tollefsen (2004) groups can be epistemic agents if they are

²³The latter requirement is less stringent than it may seem. All that is required is that the first steps toward a solution are knowable. This could be fulfilled even by knowing that one can form a committee to explore potential solutions.

²⁴Intentional epistemic agent is the one “whose wants or desires cause (at least in some significantly contributory sense) its belief formation in regards to p at t_n .” (Corlett, 2007, p. 238).

engaged in a joint rational deliberation. Moreover, collective agency has been emphasized as particularly relevant in the context of scientific inquiry (e.g. Fagan, 2012; Fagan, 2014; Gilbert, 2000; Miller, 2015a; Rolin, 2008; Wray, 2007). In view of this, assigning the norm of epistemic responsibility to an unorganized collective of epistemic agents, capable of preventing a pending epistemic harm, should not be controversial.²⁵

Furthermore, although the strict version of the agency condition must be denied by the proponent of EDJF, our account does not entail that collectives with no organization whatsoever can be the bearers of irreducible responsibility. EDJF's D1 responsibilities are reducible: the unorganized group only has a collective responsibility derivatively, in virtue of each group member's individual duty to communicate and enlist others. Meanwhile, the D2 responsibilities are only active once a suitable number of group members have joined forces and formed a mobilized group. As discussed, mobilized groups have a distinct status from the previously unorganized group. However, mobilization does not rise to the level of group agency.²⁶ Note that this is also why our theory is plausible in cases like *Abandoned-research*. The group of members who have newly joined forces need not have the kind of general sensitivity or lasting organization required for agency. Nonetheless, this mobilization is clearly a distinct level of organization from the initial random collective. Thus, the irreducibly collective responsibility of D2 is born by a group that is not entirely unorganized. The (D2) duty to save hypothesis H_1 is an obligation not of scientists who just happen to be in the field; rather, it is a duty of a mobilized group whose members have agreed to join forces.

Theoretical virtues. The final set of desiderata we need to consider are theoretical virtues commonly associated with normative epistemological accounts. Let's start from parsimony. The epistemic duty to join forces does require two kinds of duties to explain the relevant cases. One of these is irreducibly collective: D2 duties are duties of the group to prevent the harm, and cannot be reduced to individual duties. These two features do make the view less simple. Still, the theory only requires kinds of responsibilities that we have independent reason to admit. D1 duties look like ordinary obligations of individuals, and are not particularly exotic. The D2 obligations are similar to the kind that have been suggested independently for other kinds of organized groups such as corporations and research labs. As we have noted, the duty is also similar to views about collective responsibility of persistent social groups. Moreover, EDJF avoids any need for appeal to entirely new kinds of reasoning structures,

²⁵Social and feminist epistemologists have long recognized the importance to epistemology of evaluating non-agents. Feminist philosophers such as Longino (2002; 1990) and Solomon (2001) have argued that we must evaluate the social structure of scientific fields prior to evaluating any individual subject's justification for accepting scientific theories. Social epistemologists like Goldman (1999), Kitcher (1990), and Strevens (2003) have highlighted the importance of evaluating scientific systems and institutions. Thus, there is significant established precedent for epistemic evaluations of non-agents.

²⁶This distinguishes the current view from an epistemic analog of Collins' "duty to incorporate", which requires a higher degree of group organization (2013).

and to irreducible obligations to completely unorganized groups (as is required for joint/shared views such as Schwenkenbecher’s (2018)). Given the view’s explanatory power in a range of cases, as illustrated above, any loss of parsimony seems minimal.

The epistemic duty to join forces meets the main desiderata of a theory of collective responsibility. The only real concern is the agency condition, and there is reason to doubt that condition in its original form.

4 Epistemic Character of EDJF

So far, we have presented cases which intuitively support the contention that there is a kind of responsibility which applies to unorganized groups in scientific contexts. We then offered a theory which accounts for this kind of collective responsibility, inspired by one from the moral domain. Moreover, we have argued that this account meets the accepted desiderata for a theory of collective responsibility. However, even if all of this is convincing, one significant issue remains unaddressed: why think the duty is genuinely *epistemic*? Why do we need a separate notion of collective epistemic responsibility to go along with the more general moral notion? Moreover, past accounts of epistemic responsibility have focused on doxastic responsibility, i.e., subjects’ responsibility for their beliefs. Most have been concerned with using the notion of positive responsibility to explain epistemic justification.²⁷ The kind of responsibility highlighted in this paper seems different. So why think it is epistemic?

4.1 EDJF and epistemic aims

Our first argument for the genuinely epistemic character of EDJF concerns the aims²⁸ underlying this kind of responsibility. Recall our initial examples. *Abandoned-research* and *Biased* are cases where, intuitively, the unorganized group in question is responsible for failing to prevent an *epistemic* harm. In both cases, scientific inquiry has been (or will be) impeded. The harm to be prevented is epistemic insofar as it is a harm to inquiry, to the pursuit of scientific knowledge or progress.²⁹ For this reason, the obligations in these cases are obligations to prevent epistemic harms.

²⁷See, for example Code (1987), Hieronymi (2008), Kornblith (1983), Montmarquet (1993), Rolin (2008), Williams (2001), and Zagzebski (1996). Miller (2015) recognizes the importance of collective epistemic action. However, he argues that the kind of responsibility one has for epistemic actions is in fact ordinary moral responsibility.

²⁸For a discussion on different characterizations of epistemic normativity see Robitzsch (2019), Chapter 3.3. Also see fn. 14.

²⁹There is an enormous variety of accounts of scientific progress. The goal of science might be the accumulation of true beliefs, knowledge, explanations (Goldman, 1999). The goal might be to increase problem solving power (Kuhn, 1962; Laudan, 1977), or it might merely be pursuit of useful tools for prediction (Duhem, 1954), or the development of theories offering empirically adequate generalizations (Fraassen, 1980). For an overview of this literature, see Chakravartty (2017). Our theory of epistemic responsibility is neutral about which of these is the proper account of the goals of science. However, whatever the goal is, it seems fair to label it epistemic: science is a paradigmatic example of epistemic activity.

Moreover, the practice of holding the groups responsible in these cases also functions to promote epistemic ends. Preventing collective epistemic harms is epistemically beneficial generally. So, the general practice of holding groups accountable for failures to do so is itself aimed at epistemic ends. What is more, EDJF allows for the assignment of epistemic blame and praise in contexts including collective epistemic harms. In particular, our account provides both the duties in view of which epistemic blame and praise can be assigned to the relevant individuals, as well as conditions which have to be fulfilled for such duties to be triggered. For example, in *Abandoned-research* we can hold the individuals comprising the given scientific community as blame-worthy (assuming the epistemic, causal and no-defeaters conditions were satisfied).

Altogether, the epistemic character of the aims of EDJF goes a long way toward establishing that it describes a kind of epistemic responsibility. If one is satisfied by this, then nothing further is required to establish its epistemic nature.

However, some might still worry that having epistemic aims is not enough. There may be cases where an agent has an obligation to promote epistemic ends, but where this obligation is not itself epistemic. For instance, an obligation to avoid dishonesty promotes epistemic ends, but it might count as a purely moral obligation.

This thought, that epistemic aims are not enough to establish that a duty is epistemic, may be motivated by a background picture that ties epistemic obligations, duties, and reasons exclusively to belief and evidence. On this view, epistemic reasons just are reasons to believe a particular proposition is true. For a duty or obligation to count as epistemic, it must be a duty to respond to evidence with an appropriate belief state. (Feldman, 2002, Shah, 2006; Shah and Velleman, 2005). Such a view might admit that true beliefs are epistemic goods, while denying that one has an epistemic obligation, or even an epistemic reason, to pursue a greater number of true beliefs. Moreover, and more importantly for our purposes, reasons for action simply cannot be epistemic reasons, by definition. For similar reasons, obligations and responsibilities to promote epistemic ends by acting in various ways will also fail to be epistemic.

The kind of restricted view of epistemic normativity just sketched should be resisted. It excludes from the purview of epistemic evaluation, and epistemology, the greater portion of our epistemic pursuits. Inquiry, most especially scientific inquiry, is a paradigmatic example of epistemic activity. The evaluation of inquiry requires more than obligations to respect evidence.³⁰ To treat the norms of proper inquiry as a matter only of moral or pragmatic interest is to ignore important epistemic concerns. Moreover, the norms of proper inquiry can come apart from obligations to respect evidence for a particular theory. A scientist might have good reason to work on a theory based on its potential fruitfulness, despite having strong evidence against it. *Abandoned-research* is

³⁰See Robitzsch, 2019 for an encompassing discussion on this issue. Robitzsch proposes a rule-consequentialist account of ‘intellectual norms’ as epistemic norms, which have belief-influencing actions and omissions as their objects, and which are such that they guide the exercise of indirect doxastic control and govern doxastic responsibility assessments.

meant to illustrate just this kind of case.³¹ Thus, some epistemic obligations and responsibilities will go beyond responsibilities to believe according to one's evidence. This is compatible with thinking there is something special about reasons to believe a proposition is true. But they are not the only kinds of epistemic considerations.³²

4.2 EDJF and epistemic practice

The above response to the worry that EDJF is not epistemic still leaves us with the question of how to distinguish practical and moral considerations from those in our now expanded notion of the epistemic. What else, beyond an epistemic aim, is needed to make a responsibility epistemic? We suggest that one way to specify such responsibilities is as those that are embedded in a historical epistemic practice. In the cases at issue, the obligations that are intuitively present, and which are described by EDJF, exist partially in virtue of the scientific practice the researchers are engaging in. If the individuals in *Abandoned-research* were not members of their research community, engaged in the practice of medical research, they would bear no responsibility for ensuring that H_1 is pursued. Their obligations in this case are derived from their more general commitments as scientists in this field and community. By joining the scientific community, one displays what Kutz calls a 'participatory intention': "an intention to do my part of a collective act, where my part is defined as the task I ought to perform if we are to be successful in realizing a shared goal." (Kutz, 2000, p. 10). Thus, the engagement of the scientists in our example with the practice of medical research grounds our ability to normatively assess them from the standpoint of epistemic responsibility.

The basic idea is that scientists are engaged in a practice of inquiry.³³ A practice is an area of human activity with a particular set of aims, norms, and traditions. The norms of a practice are internal to that practice: a practitioner's performance can be evaluated according to these norms in a way that is independent of other forms of evaluation. Examples of practices are things like chess, basketball, and archery. These help illustrate the independence of evaluation: One can be an excellent archer while being terribly impolite, or morally bad. A chess move can be excellent according to the norms of chess, but morally repugnant because it is meant to embarrass a clearly outmatched opponent.

³¹This point has long been accepted in the philosophy of science and social epistemology. The historicist literature about theory change and pursuitworthiness in philosophy of science focuses on these facets of epistemic normativity, as do the social epistemologists concerned with social structure and distribution of labor. For an overview of historicism, see Nickles (2017). For discussion of the social epistemology, see Goldman (1999). As we will mention below, the literature on epistemic injustice, especially non-testimonial injustice, also shows a sensitivity to this (McKinnon, 2016).

³²See Steel (2010), Sosa (2019), and Fleisher (2018) for accounts distinguishing different kinds of epistemic considerations. For other projects which expand the scope of epistemic evaluation, see Singer and Aronowitz (forthcoming) and Friedman (forthcoming).

³³The notion of practice we are appealing to here is inspired by the theories of MacIntyre (1983), Sosa (2015), Longino (1990), and Solomon (2001).

Fields of inquiry can be understood as practices in this sense. The scientists in *Abandoned-research* are taking part in a particular practice of inquiry, one which has its own aims, norms, and traditions. The tradition and aims help determine the norms and standards by which research activity within the field is evaluated. These norms and standards govern when performances made in the practice count as legitimate or appropriate ways of pursuing the goals of the practice. Whether or not an individual or a group has an obligation in a particular case will therefore depend on those norms and standards of their field of inquiry. In *Abandoned-research*, the medical researchers have an obligation to continue work on H_1 . Failing to do so will mean failing to abide by the norms of the field regarding how the field's aim is pursued. These norms are epistemic norms, because the practice they are embedded in is an epistemic one, which pursues epistemic goals. This thereby provides another reason to think the obligations here are epistemic.

4.3 Epistemic responsibility

The preceding discussion suggests the following picture of what is sufficient for something to count as an epistemic responsibility:

Epistemic Responsibility: A responsibility R is epistemic if:

1. it functions to promote an epistemic aim, and
2. it takes places in the context of an epistemic practice, and is thereby appropriately evaluable by the epistemic standards of that practice.³⁴

This sufficient condition clearly makes the epistemic duty to join forces an account of epistemic responsibility. The resulting picture actually makes EDJF fit better with previous accounts of epistemic responsibility than one might have expected.

As we noted above, most theories of epistemic responsibility have been concerned with doxastic responsibility, generally in the pursuit of a theory of doxastic justification. In addition, social epistemologists have developed accounts of collective epistemic responsibility addressing the question: under what conditions does a group of agents responsibly form and hold a given knowledge claim (e.g. Corlett, 2001; Rolin, 2008; Rolin, 2016). We can dub both types of previous approaches *positive* accounts of epistemic responsibility.

In contrast, the question addressed in this paper belongs to what we call *preventionist* accounts of epistemic responsibility. While in some cases norms proposed by the two accounts may overlap,³⁵ in a number of others previous

³⁴The practice of a particular scientific field, as described above, is one kind of epistemic practice. But plausibly there are other more general kinds of epistemic practices, e.g., the practice of giving and receiving testimony, which may be developed through either biological or cultural evolution. This point is important for applications of EDJF outside of scientific inquiry. Note also that this condition requires that the responsibility in question promotes an epistemic aim. This rules out responsibilities, and practices, which promote goals or aims of ignorance or epistemic injustice.

³⁵For example, we could use Rolin's account of collective epistemic responsibility to argue that in order to avoid the harm of reaching a consensus on a false claim, the given scientific

accounts won't be up to the task. The reason for this is that the cases motivating our approach concern situations where no specific knowledge claim may be irresponsibly held, but where an epistemic harm may nonetheless occur.

Finally, it is worth mentioning that both positive and preventionist accounts of collective epistemic responsibility are, generally speaking, compatible with theories of strictly doxastic responsibility. The latter are generally concerned with belief-formation, and hold that we are generally directly responsible for what we do, and only indirectly responsible for states of affairs. Since beliefs are states, while belief-formations are actions, traditional responsibilist theories of epistemic justification (and epistemic virtue) also concern responsibilities for actions. The primary difference between such accounts and those addressing collective epistemic responsibility is that the standards of belief formation an agent is responsible for following are more constrained. That is, responsibly formed beliefs really should only be sensitive to considerations about whether the proposition in question is true.³⁶

5 Outlook and conclusion

We have argued that there are cases of collective epistemic harm, which seem to involve collective responsibility. We then offered an account of collective epistemic responsibility, the epistemic duty to join forces, which explained how to make sense of this kind of responsibility. As such, our account is a 'preventionist' one, complementing previously proposed 'positive' accounts of collective epistemic responsibility.

We conclude the paper by highlighting the significance of our account for discussions on both collective epistemic and collective moral harms. First, the success of EDJF, provides support for the fruitfulness of the moral duty to join forces as a theory of collective responsibility and it supplements it in relevant ways. For instance, our notion of a mobilized group may be helpful in the analysis of moral harms in addition to the epistemic ones. Furthermore, the EDJF offers a successful account of epistemic responsibility that also appeals to a weakened agency condition, which suggests that this revision is generally appropriate.

Second, EDJF sheds light on the importance of studying the process of mobilization, which has so far been neglected by philosophers of science, but which may be necessary for the prevention of certain epistemic harms. In particular, socio-epistemic conditions which could incentivize individuals to mobilize and thereby prevent a given epistemic harm, have been largely unexplored and deserve more attention.

community should be jointly committed to defend their claims in case they are appropriately challenged (see Rolin, 2008, p. 121). This seems similar to the duties imposed by EDJF.

³⁶Many have argued that beliefs should only be sensitive to truth (Goldman, 1979b, 1999), accuracy (Joyce, 1998, 2009; Pettigrew, 2016), or evidence (Shah, 2006; Shah and Velleman, 2005). For defense of the idea that belief-formation should not be rationally sensitive to considerations of long-run inquiry, see Berker (2013a,b), Firth (1981), Greaves (2013), and Jenkins (2007).

Finally, our account can be useful in the analysis of some additional cases, beyond the context of scientific research. For instance, EDJF offers a helpful supplement to theories of epistemic injustice³⁷ since some cases of epistemic injustice will count as collective epistemic harms. For example, hermeneutical injustices that result from lack of conceptual resources (e.g. the lack of the concept ‘sexual harassment’ prior to its introduction in the 1970s) are collective epistemic harms: they harm individuals in their capacity to know, and they are only preventable by collective action. Similarly, our account can help in addressing epistemic duties arising from the threat of ‘fake news’ typical for epistemically pernicious groups, such as epistemic bubbles and echo chambers (Boyd, 2019). While positive accounts of epistemic responsibility provide guidelines for how to recognize one may be involved in an epistemically pernicious group (e.g. by evaluating the degree to which the given group has been critical, open-minded and self-reflective, Corlett, 2001), our account specifies which epistemic duties one has upon recognizing such an involvement.

Acknowledgments We are grateful to the members of the Philosophy & Ethics Group at TU Eindhoven and to the audience at the Nature of Inquiry conference at Agnes Scott College, for valuable comments on previous versions of this paper (especially to Daniel Friedman who provided a commentary on our paper). Thanks, also, to Ruth Groff and Megan Feeney.

References

- Baehr, Jason (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford University Press.
- Berker, Selim (2013a). “Epistemic teleology and the separateness of propositions”. In: *Philosophical Review* 122.3, pp. 337–393. DOI: 10.1215/00318108-2087645.
- (2013b). “The rejection of epistemic consequentialism”. In: *Philosophical Issues* 23.1, pp. 363–387.
- Björnsson, Gunnar (2014). “Essentially shared obligations”. In: *Midwest studies in philosophy* 38.1, pp. 103–120.
- Björnsson, Gunnar (forthcoming). “Collective responsibility and collective obligations without collective moral agents”. In: *Handbook of Collective Responsibility*. Ed. by Saba Bazargan-Forward and Deborah Tollefsen. Routledge.
- Boxill, Bernard (2016). “Black Reparations”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2016. Metaphysics Research Lab, Stanford University.
- Boyd, Kenneth (2019). “Epistemically pernicious groups and the groupstrapping problem”. In: *Social Epistemology* 33.1, pp. 61–73.
- Chakravartty, Anjan (2017). “Scientific Realism”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.

³⁷For an overview of the literature, see fn. 15.

- Code, Lorraine (1987). *Epistemic responsibility*. Published for Brown University Press by University Press of New England.
- Collins, Stephanie (2013). “Collectives’ duties and collectivisation duties”. In: *Australasian Journal of Philosophy* 91.2, pp. 231–248.
- Corlett, J Angelo (2001). “Collective Responsibility”. In: *Wiley Encyclopedia of Management*.
- (2007). “Analyzing social knowledge”. In: *Social Epistemology* 21.3, pp. 231–247.
- Darby, Derrick and Nyla R. Branscombe (2014). “Beyond the sins of the fathers: Responsibility for inequality”. In: *Midwest Studies in Philosophy* 38.1, pp. 121–137.
- Dotson, Kristie (2011). “Tracking epistemic violence, tracking practices of silencing”. In: *Hypatia* 26.2, pp. 236–257.
- Duhem, Pierre Maurice Marie (1954). *The aim and structure of physical theory*. Princeton: Princeton University Press.
- Elgin, Catherine (2017). *True Enough*. Cambridge: MIT Press.
- Fagan, Melinda (2012). “Collective scientific knowledge”. In: *Philosophy Compass* 7.12, pp. 821–831.
- Fagan, Melinda Bonnie (2014). “Do groups have scientific knowledge?” In: *From individual to collective intentionality: New essays*. Oxford University Press, pp. 163–186.
- Feinberg, Joel (1968). “Collective responsibility”. In: *Journal of Philosophy* 65.21, pp. 674–688.
- Feldman, R. (2002). “Epistemological Duties”. In: *The Oxford Handbook of Epistemology*. Ed. by P. K. Moser. Oxford: Oxford University Press, pp. 362–384.
- Firth, Roderick (1981). “Epistemic merit, intrinsic and instrumental”. In: *Proceedings and Addresses of the American Philosophical Association* 55.1, pp. 5–23. DOI: 10.2307/3131397.
- Fischer, John Martin and Mark Ravizza (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Fleisher, Will (2018). “Rational endorsement”. In: *Philosophical Studies* 175.10, pp. 2649–2675.
- Fraassen, B.C. van (1980). *The scientific image*. Clarendon Library of Logic and Philosophy. Clarendon Press. ISBN: 9780198244271.
- Fricke, Miranda (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Friedman, Jane (2017). “Why Suspend Judging?” In: *Noûs* 51.2, pp. 302–326. DOI: 10.1111/nous.12137.
- (forthcoming). “The epistemic and the zetetic”. In: *Philosophical Review*.
- Gilbert, Margaret (2000). *Sociality and responsibility: New essays in plural subject theory*. Rowman & Littlefield.
- Goldman, Alvin I (1979a). “What is justified belief?” In: *Justification and knowledge*. Springer, pp. 1–23.
- (1979b). “What is justified belief?” In: *Justification and knowledge*. Springer, pp. 1–23.
- (1986). *Epistemology and Cognition*. Harvard University Press.

- Goldman, Alvin I (1999). *Knowledge in a social world*. Oxford: Clarendon Press.
- Greaves, Hilary (2013). “Epistemic decision theory”. In: *Mind* 122.488, pp. 915–952. DOI: 10.1093/mind/fzt090.
- Held, Virginia (1970). “Can a random collection of individuals be morally responsible?” In: *The Journal of Philosophy* 67.17, pp. 471–481.
- Hieronymi, Pamela (2008). “Responsibility for believing”. In: *Synthese* 161.3, pp. 357–373.
- Hindriks, Frank (2018). “The Duty to Join Forces: When Individuals Lack Control”. In: *The Monist* Forthcoming.
- Hookway, Christopher (2008). “Questions, Epistemology, and Inquiries”. In: *Grazer Philosophische Studien* 77.1, pp. 1–21. DOI: 10.1163/18756735-90000841.
- Jenkins, C.S.I (2007). “Entitlement and Rationality”. In: *Synthese* 157.1, pp. 25–45. DOI: 10.1007/s11229-006-0012-2.
- Joyce, James (1998). “A nonpragmatic vindication of probabilism”. In: *Philosophy of Science* 65.4, pp. 575–603. DOI: 10.1086/392661.
- (2009). “Accuracy and coherence: Prospects for an alethic epistemology of partial belief”. In: *Degrees of Belief*. Ed. by Franz Huber and Christoph Schmidt-Petri. Synthese, pp. 263–297.
- Khalifa, Kareem (forthcoming). “Understanding, Truth, and Epistemic Goals”. In: *Philosophy of Science*.
- Kidd, Ian James, José Medina, and Gaile Pohlhaus, eds. (2017). *The Routledge Handbook of Epistemic Injustice*. Routledge.
- Kitcher, Philip (1990). “The division of cognitive labor”. In: *Journal of Philosophy* 87.1, pp. 5–22.
- Kornblith, Hilary (1983). “Justified Belief and Epistemically Responsible Action”. In: *Philosophical Review* 92.1, pp. 33–48.
- Kuhn, Thomas S. (1962). *The structure of scientific revolutions*. University of Chicago Press.
- Kutz, Christopher (2000). “Acting together”. In: *Philosophical and Phenomenological Research*, pp. 1–31.
- Kvanvig, Jonathan L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Laudan, Larry (1977). *Progress and its problems: Toward a theory of scientific growth*. University of California Press.
- List, Christian and Philip Pettit (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press.
- Longino, Helen (2002). *The Fate of Knowledge*. Princeton: Princeton University Press.
- Longino, Helen E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, New Jersey: Princeton University Press.
- MacIntyre, Alasdair C. (1983). *After virtue: A study in moral theory*. University of Notre Dame Press.
- McGary, Howard (1986). “Morality and Collective Liability”. In: *Journal of Value Inquiry* 20.2, pp. 157–165.

- McKinnon, Rachel (2016). “Epistemic injustice”. In: *Philosophy Compass* 11.8, pp. 437–446.
- Medina, José (2013). *The epistemology of resistance: Gender and racial oppression, epistemic injustice, and the social imagination*. Oxford University Press.
- Millar, Boyd (2020). “Shared epistemic responsibility”. In: *Episteme*, pp. 1–14.
- Miller, Boaz (2015a). “Why (some) knowledge is the property of a community and possibly none of its members”. In: *The Philosophical Quarterly* 65.260, pp. 417–441.
- Miller, Boaz and Isaac Record (2013). “Justified belief in a digital age: On the epistemic implications of secret Internet technologies”. In: *Episteme* 10.2, pp. 117–134.
- Miller, Seumas (2015b). “Joint epistemic action and collective moral responsibility”. In: *Social Epistemology* 29.3, pp. 280–302.
- Millson, Jared and Kareem Khalifa (2020). “Perspectives, Questions, and Epistemic Value”. In: *Knowledge From a Human Point of View*. Ed. by Michela Massimi. Springer Verlag.
- Montmarquet, James A. (1993). *Epistemic Virtue and Doxastic Responsibility*. Rowman & Littlefield.
- Nickles, Thomas (2017). “Historicist Theories of Scientific Rationality”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.
- Parfit, Derek (1984). *Reasons and persons*. Oxford University Press.
- Pettigrew, Richard (2016). *Accuracy and the laws of credence*. Oxford University Press UK.
- Pettit, Philip and David Schweikard (2006). “Joint actions and group agents”. In: *Philosophy of the Social Sciences* 36.1, pp. 18–39.
- Pinkert, Felix (2014). “What we together can do”. In: *Midwest Studies in Philosophy* 38.1, pp. 187–202.
- Robitzsch, Andrea (2019). *An Externalist Approach to Epistemic Responsibility: Intellectual Norms and Their Application to Epistemic Peer Disagreement*. Vol. 411. Springer.
- Rolin, Kristina (2008). “Science as collective knowledge”. In: *Cognitive Systems Research* 9.1-2, pp. 115–124.
- Rolin, Kristina Helena (2016). “Collective epistemic responsibility: A reply to Chris Dragos”. In: *Social Epistemology Review and Reply Collective* 5.11, pp. 7–11.
- Schwenkenbecher, Anne (2018). “Making sense of collective moral obligations: A comparison of existing approaches”. In: *Collectivity: Ontology, Ethics, and Social Justice*. Ed. by Kendy M Hess, Violetta Igneski, and Tracy Isaacs. Rowman & Littlefield, pp. 109–132.
- (2019). “Collective moral obligations: ‘we-reasoning’ and the perspective of the deliberating agent”. In: *The Monist* 102.2, pp. 151–171.
- Šešelja, Dunja, Laszlo Kosolosky, and Christian Straßer (2012). “Rationality of Scientific Reasoning in the Context of Pursuit: drawing appropriate distinctions”. In: *Philosophica* 86, pp. 51–82.

- Šešelja, Dunja and Christian Straßer (2014a). “Epistemic justification in the context of pursuit: a coherentist approach”. In: *Synthese* 191.13, pp. 3111–3141.
- (2014b). “Heuristic Reevaluation of the Bacterial Hypothesis of Peptic Ulcer Disease in the 1950s”. In: *Acta Biotheoretica* 62, 429–454.
- Shah, Nishi (2006). “A new argument for evidentialism”. In: *Philosophical Quarterly* 56.225, pp. 481–498. DOI: 10.1111/j.1467-9213.2006.454.x.
- Shah, Nishi and J. David Velleman (2005). “Doxastic deliberation”. In: *Philosophical Review* 114.4, pp. 497–534. DOI: 10.1215/00318108-114-4-497.
- Singer, Daniel J. and Sara Aronowitz (forthcoming). “What Epistemic Reasons Are For: Against the Belief-Sandwich Distinction”. In: *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*. Ed. by Billy Dunaway and David Plunkett.
- Smiley, Marion (2017). “Collective responsibility”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.
- Solomon, Miriam (2001). *Social Empiricism*. Cambridge, Massachusetts: MIT press.
- Sosa, Ernest (2015). *Judgment & agency*. Oxford University Press UK.
- (2019). “Suspension as Spandrel”. In: *Episteme* forthcoming, pp. 1–12.
- Steel, Daniel (2010). “Epistemic values and the argument From inductive risk”. In: *Philosophy of Science* 77.1, pp. 14–34. DOI: 10.1086/650206.
- Strevens, Michael (2003). “The role of the priority rule in science”. In: *The Journal of philosophy* 100.2, pp. 55–79.
- Tollefsen, Deborah (2004). “Collective epistemic agency”. In: *Southwest Philosophy Review* 20.1, pp. 55–66.
- (2015). *Groups as agents*. Polity.
- Williams, Michael (2001). *Problems of knowledge: A critical introduction to epistemology*. Oxford University Press.
- (2008). “Responsibility and reliability”. In: *Philosophical Papers* 37.1, pp. 1–26.
- Williamson, Timothy (2002). *Knowledge and its Limits*. Oxford University Press on Demand.
- Wray, K Brad (2007). “Who has scientific knowledge?” In: *Social Epistemology* 21.3, pp. 337–347.
- Wringe, Bill (2016). “Collective obligations: Their existence, their explanatory power, and their supervenience on the obligations of individuals”. In: *European Journal of Philosophy* 24.2, pp. 472–497.
- (2019). “Global obligations, collective capacities, and ‘ought implies can’”. In: *Philosophical Studies*, pp. 1–16.
- Zagzebski, Linda Trinkaus (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.