

How We Fail to Know: Group-Based Ignorance and Collective Epistemic Obligations

Political Studies

1–18

© The Author(s) 2021

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/00323217211000926

journals.sagepub.com/home/psx**Anne Schwenkenbecher** 

Abstract

Humans are prone to producing morally suboptimal and even disastrous outcomes out of ignorance. Ignorance is generally thought to excuse agents from wrongdoing, but little attention has been paid to group-based ignorance as the reason for some of our collective failings. I distinguish between different types of first-order and higher order group-based ignorance and examine how these can variously lead to problematic inaction. I will make two suggestions regarding our epistemic obligations vis-a-vis collective (in)action problems: (1) that our epistemic obligations concern not just our own knowledge and beliefs but those of others, too and (2) that our epistemic obligations can be held collectively where the epistemic tasks cannot be performed by individuals acting in isolation, for example, when we are required to produce joint epistemic goods.

Keywords

collective action, group knowledge, social epistemology, inaction, collective ignorance, epistemic obligations, joint epistemic goods

Accepted: 16 February 2021

Introduction

Humans are prone to producing morally suboptimal and even disastrous outcomes out of ignorance. Examples of collectively caused problems that are due to group ignorance abound: they can arise in very large populations that are not group agents in the strict sense (List and Pettit, 2011) but also in tightly structured organizations. Groups of people panic-buying essential supplies because of a pending pandemic give rise to a collective action problem that may be the result of ignorance. Suboptimal communication structures in organizations can mean partial and asymmetric levels of knowledge and ignorance, leading to members making underinformed or misinformed decisions. Another example

College of Arts, Business, Law and Social Sciences, Murdoch University, Murdoch, WA, Australia

Corresponding author:

Anne Schwenkenbecher, College of Arts, Business, Law and Social Sciences, Murdoch University,
90 South St., Murdoch, WA 6150, Australia.

Email: A.Schwenkenbecher@murdoch.edu.au

is pluralistic ignorance, which can be responsible for publicly upholding social norms that people privately no longer agree with (Bicchieri, 2017). It is often the case that we fail to do good because we do not know *what* to do or even *that* we are able to produce some good at all. Sometimes, we may be blamed for collectively failing to do good or to prevent harm, but when?

Ignorance is generally thought to excuse agents from wrongdoing: if we are blamelessly ignorant of some feature of our action or unaware that some consequence would follow from it we will usually not be blamed for its negative outcome. But when is ignorance blameless? According to Gideon Rosen (2004), ignorance is only blameworthy if an agent has previously failed in fulfilling their epistemic obligations. Daniel Miller (2017: 1568) recaps Rosen's position as follows: 'epistemic obligations amount to obligations to do certain things that will or might result in an improved epistemic position with respect to one thing or another'.

Whereas *individual* agents' ignorance and its impact on their moral responsibility and blameworthiness have attracted a fair amount of attention in the literature, hardly anything has been written on ignorance as it can obtain in groups of agents, even though the phenomenon described seems ubiquitous. One of the few exceptions is Såde Hormio (2018: 7), who – fittingly – remarks that 'literature on collective ignorance and what it means for responsibility has so far been quite thin on the ground'.¹

Having said that, philosophers who have written about responsibility for collective inaction tend to invoke an epistemic condition. Very roughly, it goes like this: agents can be held responsible for failing to collectively produce an outcome only if there was an obvious or salient solution to a collective action problem that a reasonable person should have been aware of (Held, 1970; Isaacs, 2011; May, 1992; Petersson, 2008; Pinkert, 2014). It is fair to say that these philosophers do not specifically address the question of what *type* of knowledge must obtain within a group in order to successfully address joint necessity cases and, relatedly, what types of ignorance would stymie collective endeavours in a way that removes blameworthiness.

This article attempts to close both of these gaps in theorizing. It explores group-based, or collective, types of ignorance and, correspondingly, examines what our collective epistemic obligations vis-a-vis collective (in)action problems are. Before I continue, let me add a caveat: I am using the term 'groups' fairly loosely here to cover all collectivities from so-called 'random collections' – pluralities of agents that are in principle capable of intentionally producing certain outcomes or performing actions together such as arbitrary bystanders (Held, 1970) – to group agents such as corporations (List and Pettit, 2011). The differences between these groups matter for my main argument only when it comes to collective epistemic obligations, which is where I will directly address this issue.

In the next section ('Group-Based Ignorance'), I discuss varieties of group-based ignorance. Then, I will demonstrate how group-based ignorance can lead to morally suboptimal collective outcomes ('How Collective Ignorance Can Affect the Collective Production of Morally Important Goods'). Finally, I will turn to the question of what it means for groups to have collective epistemic obligations ('Collective Epistemic Obligations'). Collective inaction resulting from group-based ignorance is blameworthy where collective epistemic obligations have been violated.

Group-Based Ignorance

Some of our collective failings are due at least in part to ignorance. What exactly does it mean to be ignorant of something? According to the *standard view* defended by Pierre Le

Morvan (2011), ignorance is the lack of knowledge. In other words, an agent x is ignorant of a true proposition p if she does not know that p is true. In contrast, Rik Peels (2012) argues that ignorance is the lack of true belief. Epistemologists do not always agree on how ‘knowledge’ is best understood. For the purpose of this article I will use it to mean ‘true belief’ and as such my definition of ignorance will follow Peels’ *new view*. Therefore, when I say that an agent is ignorant of (a true proposition) p , I mean that she does not hold a true belief with regard to p . She could either hold a false belief or no belief or suspend belief (agnosticism).²

Le Morvan further distinguishes between *propositional* and *factive* ignorance (Peels, 2012):

X is propositionally ignorant with regard to p if she does not know of the proposition p .³

X is factively ignorant if she does know of p , but does not know that p is true.⁴

With regard to climate change, for instance, we can say that, largely, people were propositionally ignorant of it until late in the twentieth century. In other words, they were ignorant in that they were not holding any beliefs concerning anthropogenic climate change; they were not even aware that there was anything to know.⁵ In contrast, at least some of those who *now* deny that anthropogenic climate change is occurring are factively ignorant: they know of the proposition ‘Human activity causally contributes to climatic change’, but they mistakenly believe that this proposition is false or that the facts of the matter are not sufficiently settled. Similarly, those who believe that vaccinations cause autism and therefore do not vaccinate their children are factively ignorant with regard to the proposition ‘Vaccination does not cause autism’. By vaccinating ourselves and our children we are not only acting in our own interest, but also, through herd immunity, protect those vulnerable people who for one reason or another cannot be vaccinated. Those unaware of the fact that a drop in vaccination rates can extinguish herd immunity and cause people to die who would have been otherwise protected are *propositionally* ignorant of the harm they are causing in conjunction with others who likewise reject vaccines.

In contrast to the singular way in which individuals can know (or be ignorant of) some proposition p , there are many different ways in which sets of individuals can know or be ignorant of some proposition p . In other words, there are many different types of group-based ignorance. I will begin by explaining group-based knowledge and then move on to group-based ignorance.

Types of group knowledge differ along at least two dimensions: interconnectedness and symmetry. *Interconnectedness* refers to the degree to which individual group members know what other group members know. The level of interconnectedness is greatest where there is common knowledge and lowest where knowledge is shared, that is, where a set of individuals all have the same true belief but without knowing others’ beliefs. *Symmetry* refers to the way in which knowledge is distributed, that is, whether some group members know more than others with regard to the issue at hand. Where the distribution of knowledge is asymmetric, some group members hold more information than others and may play a coordinating role (Roy and Schwenkenbecher, 2019).

Let me start with *shared knowledge* wherein knowledge is distributed but not interconnected:

A proposition p is *fully shared* [and symmetrically distributed] *knowledge* in a set of agents [a, b, \dots, n] if each of these agents holds a true belief that p . Further, we can say that a proposition

is *widely shared* amongst agents if a large subset holds a true belief that p . It is *partially shared* if a small subset of agents holds a true belief that p . Widely and partially shared knowledge come in degrees.⁶ A proposition p is no longer shared knowledge of any kind if only one group member holds a true belief that p .

The concepts of widely and partially shared knowledge may not seem pertinent where small groups containing only two agents are concerned, but they are useful when it comes to very large groups like humanity, the educated, the global rich, or the citizens of any particular state. They allow us to capture a (larger) group's saturation with a belief or else a measure of how widespread ignorance of that belief is in the group:

Correspondingly, *fully shared ignorance* of a proposition p obtains in set of agents $[a, b, \dots, n]$ if *every*⁷ agent in that set is ignorant of p .⁸ *Widely shared ignorance* of a proposition p means that a large subset of agents is ignorant of p while the remaining group members are not, with the number of ignorant group members greater than the number of group members in the know. *Partially shared ignorance* of a proposition p means that a small subset is ignorant of p while the remaining group members know p , and more people than not know p . Ignorance is shared only where more than one group member is ignorant of p .

Shared ignorance is first-order ignorance – it is a shared lack of true first-order beliefs. Further down, I will turn to higher order ignorance, which is a lack of true second- or higher order beliefs, as in where I false believe that you do not believe p .

For certain types of coordination problems to be (reliably) resolved or for some joint activities to be successfully carried out, every member of a group will need to share true beliefs concerning some proposition(s). For instance, in order for you and me to meet at the Weston Library Café, we both need to know not only which café is meant, but also where it is. Now, obviously we may end up at the same café by sheer coincidence, but this is not a robust possibility. Here, any one individual's ignorance will result in failure of the joint activity.

However, other joint activities may fail only where there are multiple ignorant agents, that is, where ignorance is shared. If my choir is scheduled to sing at the community arts festival then we need enough singers from each voice section (sopranos, altos, etc.) in order to be able to perform. That is, we need enough people to have the relevant information so they turn up at the right place at the right time, but it is not necessary that each and every member of the choir is present.

In general, which type of ignorance (and resulting inaction) prevents people from coordinating or acting together depends on the type of joint necessity case they are facing. *Strict joint necessity*⁹ cases are those where the number of agents available for addressing the problem in question (or producing the outcome in question) equals the minimum number of agents required for solving the problem (or producing the outcome). Take the above example of two people meeting at a specific cafe. Here any one individual agent's ignorance suffices for frustrating coordination.

In contrast, in *wide joint necessity* cases, an individual agent's ignorance will by itself not undermine the realization of some joint activity. We are facing *wide joint necessity* if not every available agent (or every member of the group) is required for performing the joint action in question (or producing the collective outcome). Not all agents need to act and thence be in the know, as in the example of the choir. There must be an *adequate degree* of generally shared knowledge of the relevant proposition (concerning the problem and the individual contributions to its solution). Most large-scale joint necessity

cases are wide rather than strict. There, collective outcomes can be produced even if not all available agents contribute. One example is the realization of herd immunity from infectious diseases: it does not require a 100% vaccine coverage rate in order to generate herd immunity. For instance, for measles the herd immunity threshold is 93%–95% (Funk et al., 2019). Shared ignorance can be a factor when we fail to produce public goods.

Let me now turn to higher order group-based ignorance, starting with explaining the corresponding type of group-based knowledge (or true belief). While shared knowledge and its derivatives are a type of aggregate knowledge, there also exist forms of genuinely group-based, *interconnected*, knowledge. The strongest form of such knowledge is *common knowledge*. A well-accepted definition of common knowledge is the iterative definition:

A proposition p is common knowledge in a set of individuals $[a, b]$ if it is true that a knows that p , b knows that p , and a knows that b knows that p , and b knows that a knows that p , and so on. Common knowledge of p can exist in a set of agents $[a, b]$ at different levels, depending on how many levels of iteration of ‘ a/b knows that b/a knows’ are true.¹⁰ Another way of putting it is to say: A proposition p is (level one) common knowledge between two individuals $[a, b]$ if p is shared knowledge between them *and* if both know that this is so.

This means that common knowledge of a proposition p can fail in two ways: either at the level of shared knowledge (first-order failure) or at a higher level n (n^{th} -order failure). Sometimes, people who share first-order knowledge of p may lack second-order knowledge of p (as appears to be the case for *pluralistic ignorance*, which is explained below).

Here is an example of second-order ignorance: you and I talk about going to a live music gig together that we vaguely heard of but we do not know exactly where and at what time it takes place. As we part, we agree that whoever finds out first will tell the other. Later that day you message me that it is at Isis Farmhouse at 3 pm. As I am reading the message it is now the case that each of us knows where and when the gig takes place. First-order knowledge about the place and time is now shared, but you do not know that it is shared; only I do. There is asymmetric, or partial, second-order knowledge. You will not have second-order knowledge unless I confirm having received your message. Let us suppose that for some reason I do not confirm with you. As a result your conditional intention of going to the gig *if I also go* does not become an actual intention to go to the gig, because you have no reason to believe that I will be there. My failure to confirm means that I fail to generate higher order knowledge in both, which in turn undermines our plan to meet.¹¹ This is due to partial second-order ignorance. We can see that in larger groups, second-order ignorance can come in degrees: it can be widely or partially shared.

Where second-order ignorance obtains it must be factive (rather than propositional):¹² A and B both believe that p , but at least one of them does not hold a true belief that the other person believes that p . Of course, this suggests that social or group-based ignorance is ubiquitous – after all there is an endless number of beliefs that others hold in common with us (shared knowledge) and of which we falsely assume them to be ignorant of (or where we suspend judgement).

Let me return to ‘common knowledge’, though, which is a very demanding concept in its iterative form. For the purpose of understanding large-scale collective action problems it is actually useful to have a weaker notion of higher order knowledge:

Public knowledge: A proposition p is publicly known in a group if knowledge of p is widely shared and there is widely shared knowledge concerning that very fact (Schwenkenbecher, 2021).

For example, it is widely shared knowledge that high levels of greenhouse gas emissions are causing climate change and that fact itself is widely known. I know that you (are very likely to) know this without having even met you or spoken to you about it. In other words, a proposition is publicly known in a group if most people know the proposition and that there is a general awareness that this is something people know even if not every single person does. Another example: it is public knowledge that too much sugar is bad for you. We generally assume that people know this (even if a specific person might not):

Public ignorance of a proposition p obtains in a group where p is widely shared knowledge (or a widely shared true belief) in that group, and the group members (either fully or widely) share ignorance of that fact.

Take the example of vaccination again. Let us assume that in a population it is public knowledge that vaccinating your children both protects them from contracting an infection *and* generates the important public good of herd immunity. Furthermore, a subset of that population shares the true belief that individual failures to vaccinate will not undermine the production of that public good as long as these remain – in aggregation – below the herd immunity threshold. Even if the members of that subgroup are morally motivated and conscientious, (falsely) believing that they are the only ones knowing that a certain number of defections will make no difference to the outcome (while also believing that everyone else knows how important it is to produce the good) may provide them with sufficient reason for defecting or else with insufficient reason for cooperating. Their rationale may be this: if my decision not to vaccinate has no negative impact on the production of the public good, but it has a potentially positive impact on the well-being of my child or is simply more convenient then it seems that I should defect (i.e. not vaccinate). If this ‘harmless defection belief’ is fully or widely shared, but not publicly known, and instead people widely (and falsely) believe to be alone in holding the belief herd immunity will be undermined. People may defect (fail to contribute to the public good) on the basis of this false second-order belief. The point here is not that (factive) ignorance of others’ beliefs will make people defect. Rather, it is that factive ignorance (in this case a false second-order belief) gives people a reason to defect that they otherwise would not have. In other words, public ignorance concerning the production of herd immunity and harmless defection rates can result in a failure to produce this important public good.

Let me conclude this section by briefly pointing out in what sense shared (first-order) ignorance and public (and other types of higher order) ignorance are ‘collective’ or ‘group-based’. Readers may grant that second-order (or higher order) ignorance is undeniably social, and in that sense group-based, because it concerns the interconnection between two (or more) people’s epistemic states. But readers may be less convinced that shared ignorance – the shared lack of first-order true belief – is group-based or social in a meaningful sense. After all, where a set of agents are ignorant in this sense it is really just the individual agents in that set who each lack the relevant beliefs concerning the same fact or proposition, but their epistemic states are not interconnected. Seumas Miller (2018: 28) argues that this form of ‘aggregate ignorance’ is not *collective* ignorance because it lacks interdependence and interconnection. I am not convinced that much

hangs on drawing a firm line between collective ignorance on one hand and aggregate ignorance on the other. Treating shared ignorance as a collective phenomenon allows for a focus on *groups* of agents and (the effect of) their combined epistemic states. Focusing on sets of agents rather than individuals enables us to pinpoint the social significance of ignorance as it is spread in groups and populations.

How Collective Ignorance Can Affect the Collective Production of Morally Important Goods

In this section, I will run through various examples of how higher order group-based ignorance can be detrimental to the production of morally important goods.

Pluralistic ignorance is a particularly salient example of what I have above called public ignorance. It is present where members of a reference group share the same first-order belief but they hold false second order beliefs regarding others' first-order beliefs (or expectations) (Bicchieri, 2017; Bicchieri and Mercier, 2014).¹³ This type of second-order ignorance can be a reason for upholding social norms even where privately people disagree with that norm. In addition, it can potentially explain why people fail to act in a moral emergency or on a pressing social issue.

According to Cristina Bicchieri (2017: 42), pluralistic ignorance is 'a cognitive state in which each member of a group believes her personal normative beliefs and preferences are different from those of similarly situated others, even if public behaviour is identical'. Under those circumstances, 'All end up conforming to the public norm, oblivious to the possibility that they are participants in a group dynamic in which all pretend to support the norm, while in fact all dislike it' (Bicchieri, 2017: 44). This so-called 'belief-trap' may be hard to escape, because it is not straightforward to change a social norm (Bicchieri, 2017). Having collaborated for several years with UNICEF on changing harmful practices such as female genital cutting (FGC), Bicchieri and Mercier (2014) suggest that private support of practices such as FGC is often significantly lower than the practice's prevalence and the level of support for the practice that people infer based on its prevalence. Because social norms are 'supported by shared normative beliefs', in order to change actual behaviour, we must change people's second-order beliefs about other people's normative views.¹⁴

In two studies, psychologists Nathaniel Geiger and Janet K. Swim show that pluralistic ignorance and self-silencing are regular occurrences with regard to climate action:

Survey respondents who did not themselves doubt climate change were less willing to discuss the topic when they inaccurately believed fellow students would not share their opinion than when they accurately perceived they were in the majority' and 'when accurate portrayals of others' beliefs were presented, those who were concerned about climate change were more willing to discuss the topic. (Geiger and Swim, 2016: 88)

Geiger and Swim (2016) suggest that one way to increase public engagement with climate change is 'to correct pluralistic ignorance, informing them that a majority of others share their concern'.¹⁵

Pluralistic ignorance is regularly seen as responsible for the 'bystander effect' – the kind of scenario where several people observe a moral wrong that should be remedied, but they fail to act because each assumes that the failure of others to intervene is reflective

of their inner conviction that no action is warranted. This leads them to doubt that they are legitimately worried about the scenario they are witnessing or else it makes them afraid to step out of line for fear of embarrassment (D. T. Miller and McFarland, 1987). Basically, it is a misattribution of motives to others' observed behaviour, or an illusion of uniqueness: while all behave identically, each bystander attributes to the other bystanders motives that differ from her own.

The previous examples reflect different ways in which group-based ignorance can *directly* undermine the production of collective goods or stymie joint efforts to coordinate and cooperate.¹⁶ Ignorance impacts subjective reasons for acting. Generally, first-order group-based ignorance of some proposition x will make it *at least unlikely* that people produce an collective outcome (or perform a joint action) based on x where this requires a conscious change of their behaviour.¹⁶ It is in that sense that collective inaction with regard to some joint necessity problem can be due to first-order (shared) ignorance.

Higher order ignorance can stymie collective endeavours in a several ways. For instance, public knowledge is crucial to social coordination and the efficacy of social norms. The fact that moral norms are common knowledge (we all know that we all know that killing is wrong, for instance), or at least public knowledge to a very high degree, explains some of their motivating force (Louis et al., 2004).

Apart from this direct impact of public ignorance on collective action, it may also have an indirect impact. Social psychologists have shown that beliefs in collective efficacy are crucial for motivating people to act on social issues (Thomas et al., 2009). If I believe my contribution to a collective or public good to be in vain because I falsely assume that others are ignorant of either the good's importance or of how to produce it, then it is not reasonable for me to contribute to the good. I may conclude that it is futile to do may share as well as expect that I will not be reproached for this omission. If many people hold such (false) second-order beliefs their resulting individual omissions can produce morally tragic collective inaction.

Furthermore, according to Marion Godman (2013), joint action is regularly socially motivated. Social motivations, writes Godman (2013: 590), are 'emotional and affective factors', including social emotions such as empathy. They 'help shape future behaviour by prompting a continued engagement with both the activities and individuals in joint action' by generating a 'shared perspective' (Godman, 2013: 594–595). Crucially, empathy with another person's feelings requires knowledge of those feelings (whether through explicit communication or through facial or bodily expression and observation). To the extent that empathy and other social emotions do motivate people to contribute to joint causes (directly or via the formation of social bonds), ignorance of others' emotional states will weaken the likelihood of collective endeavours. Both are examples of how ignorance can impact on people's social motivation to contribute to collective endeavours.

Another area of collective action where group-level beliefs play an essential role is in forming social identities, which in turn are a crucial element in motivating action for social and political causes. According to social psychologists Emma Thomas and Winnifred Louis (2013: 178), 'it has been shown that a common social identity can facilitate greater cooperation in social dilemmas, group-based helping . . . , and inferred trust based on the shared social relationship'. They write that '[p]eople engage in collective action because they identify with groups. A social identity acts to link individual and group' (2013: 176). Quoting fellow social psychologist Henri Tajfel, they explain: 'More formally, social identities are commonly defined as 'that part of the individual's

self-concept which derives from his or her knowledge of membership in a social group (or groups) together with the value and the emotional significance attached to it' (Tajfel, 1981: 255, quoted in Thomas and Louis 2013: 177)'. According to Thomas and Louis, '[w]hen a personal identity is salient (or psychologically operative), behavior will be defined by individual-level interactions; however when a social identity is salient behavior will be defined by group-level interactions' (2013: 177). Importantly, social identity is regularly 'opinion-based', that is, it arises from common beliefs, for instance, in the justice or injustice of certain states of affairs. According to Thomas et al.:

where people come to see themselves as a collective defined by a shared opinion, they would be expected to adhere to the norms of that group. Where those opinion-based groups are defined by opinions and norms relating to positive social change, then identification with such groups is likely to lead to behavioural change in support of prosocial actions. (Thomas et al., 2009: 199)

Thomas et al. use the example of fighting global poverty:

. . . identifying as a member of an opinion-based group that is defined by a shared movement that 'we should act to overcome poverty in developing nations' could be a plausible and useful way to promote a sustained orientation toward social action in this context. (Thomas et al., 2009: 200)

In other words, people are more likely to get behind social and political causes if they identify as members of a group who share the same beliefs concerning the importance of those causes. Note that in order to be motivated by the abovementioned beliefs, group members would have to share both first- and second-order beliefs. Group-based ignorance that undermines the forming of such identities and therewith stymies social action can exist on either level. Furthermore, Thomas et al. (2013) identify a (first order) shared belief in the group's *efficacy* in relation to such causes as crucial in motivating action.

For completeness' sake let me point to another way in which shared ignorance and public knowledge can tragically interact. Sometimes, even though people are motivated and willing to act, there is shared ignorance concerning how to solve a particular collective action problem, for instance, which individual actions will positively solve a problem or how to coordinate individual actions, and at the same time there is public knowledge of that fact. That is, we all know that (we all know that) none of us knows how to resolve that problem. This is public knowledge of shared ignorance and it is especially disheartening.

To sum up, group-based ignorance – shared (first order) or public (higher order) – may lead to collective inaction in a number of ways: by making it unlikely that people form relevant intentions to remedy a joint necessity problem in the first place (where they are ignorant of the problem or their collective ability to fix it), it can have them misjudge the group's efficacy, it may stand in the way of forming the relevant social identity, or otherwise undermine their motivation and reasons for taking action if they false judge to be the only ones concerned about an issue (as in the case of pluralistic ignorance).

Collective Epistemic Obligations

I established that ignorance can undermine collective action and the production of morally important goods in a variety of ways. Let me now return to my initial question: When are we to blame for our ignorance? Scholars seem to agree that blameless ignorance excuses wrongdoing, but disagree on when ignorance is blameless. On one view,

ignorance is only blameworthy where an agent has previously failed in fulfilling their epistemic obligations (D. J. Miller, 2017; Rosen, 2004). On this account, ‘epistemic obligations amount to obligations to *do* certain things that will or might result in an improved epistemic position with respect to one thing or another’ (Miller, 2017: 1568, on Rosen).

As mentioned before, philosophers who think that collective inaction can be morally wrong hint at epistemic conditions for responsibility and blameworthiness but do not specifically spell out what group-based epistemic obligations are. Virginia Held (1970: 476) argued that:

when the action called for in a given situation is *obvious to the reasonable person* and when the expected outcome of the action is *clearly favorable*, a random collection of individuals may be held responsible for not taking a collective action. But when the action called for is *not obvious to the reasonable person*, a random collection may not be held responsible for not performing the action in question, but, in some cases, may be held *responsible for not forming itself into an organized group capable of deciding which action to take* (my emphasis).

While Held does not use the terms ‘belief’ or ‘knowledge’, she clearly thinks that we can be held responsible for inaction when the required collective action was jointly feasible and easily knowable but also when the group fails to improve its epistemic status vis-a-vis the emergency. Tracy Isaacs and Felix Pinkert, similar to Held, invoke a reasonable person standard and ascribe obligations to groups of agents where the solution is ‘salient’ (Pinkert, 2014) or ‘clear’ and ‘in focus’ (Isaacs, 2011).¹⁷ However, it is fair to say that in general those philosophers who have explicitly written about responsibility for collective inaction do not specifically address the question of collective knowledge or a group’s obligation to address their ignorance and improve their epistemic status.¹⁸

Michael Doan criticizes Held and Isaacs for taking the so-called knowledge condition to be a necessary criterion for ascribing responsibility for collective inaction. That is, he rejects the idea that ‘with respect to any problem that demands nothing short of a coordinated response, knowledge of a clear and definitive collective action solution on the part of individuals is a prerequisite for moral responsibility’ (Doan, 2016: 537). He objects to what he perceives to be an underlying epistemological individualism, an – on his view – overly demanding epistemic standard it sets for individual knowers, and for lending support to what he calls a ‘malign form of response skepticism’: ‘the use of uncertainty and perceived unclarity as moral justification for collective inaction’ (Doan, 2016: 538). According to Doan, we share responsibility for transforming them through an ‘ongoing process of collective inquiry’, which would correct and counteract ‘individuals’ epistemic flaws’ (Doan, 2016: 544–546). Doan thinks that certain collective efforts are literally ‘unthinkable’ and that we need to ‘come to see one another as people we can learn from and as collaborators in a process of continual development’ (Doan, 2016: 551).

Doan’s positive point that the production of knowledge is itself most of the time a collective endeavour and should not (always) be seen as its precondition is plausible.¹⁹ He emphasizes the fundamentally collective nature of our epistemic obligations and the necessity of a continued improvement of our joint epistemic resources. While, unfortunately, Doan does not provide much detail on what he means by those resources (or his claim that we need to move towards creating the ‘unthinkable’), I will have a closer look at the suggestion that meeting our epistemic obligations is a collective endeavour.

In doing so, I will refine my initial question. I started out asking when ignorance (or the lack of knowledge) of a collective solution to a problem is blameless, that is, when it

does not result from a failure to meet our epistemic obligations. What exactly do I mean by *epistemic* obligations? Furthermore, in what sense can such epistemic obligations meaningfully be described as *collective*? Gideon Rosen (2003: 63) contrasts epistemic obligations that are ‘requirements, the satisfaction of which makes for epistemically justified belief’ with ‘*moral* obligations governing the epistemic aspects of deliberation. A belief may be faultless in my sense without being justified’ (Rosen, 2003: footnote 3). Likewise, I focus on epistemic obligations as *moral* obligations.²⁰ Rosen writes that, ‘We are under an array of standing obligations to inform ourselves about matters relevant to the moral permissibility of our conduct: to look around, to reflect, to seek advice, and so on’ (Rosen, 2003: 63). He argues that whether or not an action is done from moral or from factual ignorance, the agent performing it is only culpable when she is culpable for her ignorance (Rosen, 2003: 64). On Rosen’s view, ancient slaveholders may well have been blamelessly ignorant of the wrong that constitutes slavery and as such were not blameworthy for committing the morally wrong act of enslaving another human being.

He argues that we have to take epistemic precautions against negligent harm: these are procedural epistemic obligations, which are impossible to codify – ‘the person of ordinary prudence provides a serviceable heuristic’ – what would they have done in the circumstances? (Rosen, 2004: 301). He adds that ‘these procedural obligations are always obligations to *do* (or refrain from doing) certain things: to ask certain questions, to take careful notes, to stop and think, to focus one’s attention in a certain direction, etc’ (Rosen, 2004: 301). It is about taking steps ‘to ensure that when the time comes to act, one will know what one ought to know’ (Rosen, 2004: 301).²¹

For our purposes, the important question is what it means to discharge one’s epistemic obligations thus understood, focusing especially on the *social* nature of some of our knowledge (and ignorance) as well as acknowledging that improving the epistemic status of a group and its members may require collective action. Based on the preceding discussion, I will make two suggestions: (1) that our epistemic obligations concern not just our own knowledge and beliefs but those of others, too and (2) that our epistemic obligations can be held collectively where the epistemic tasks in question cannot be performed by individuals acting in isolation.

As for (1), if we believe that we sometimes have *collective* responsibility (or obligations) to prevent harm or produce a good – a claim that has been defended comprehensively elsewhere²² – then it seems that such responsibility encompasses obligations to (make an effort to) generate the requisite epistemic capacity within the respective group of agents. Generating such epistemic capacity can involve spreading information to create first-order knowledge (or preventing first-order ignorance) or communicating one’s own and others’ beliefs in order to create second- (or higher) order knowledge. Imagine a variation on Peter Singer’s (1972) shallow-pond scenario, where a joint effort of two passers-by is required to save a drowning child. It would be strange, indeed, if one passer-by were to claim that she did not rescue the child because she believed that the other one was not aware of the child drowning and so would not have assisted her. Of course, we would legitimately expect each of them to acquaint herself with the first-order beliefs of the other and to establish shared and common beliefs concerning the facts of the matter and expect both to jointly determine the best course of action.

As we have seen in several of the examples discussed in section ‘How Collective Ignorance Can Affect the Collective Production of Morally Important Goods’, including taking action on climate change or challenging outdated social norms, the success of collective endeavours often critically depends on higher order knowledge. Obviously, in

many of the above cases there will be an easy remedy to group-based ignorance; in others it will be very difficult. It is much more difficult to induce higher order knowledge in larger and dispersed groups. Yet, that is precisely what public policy announcements, media publications, information campaigns and other public broadcasts are aimed at: they ensure not only that everybody shares certain first-order beliefs, but that everybody knows that fact. In short, they generate public knowledge.

Generating shared or higher order collective knowledge in a group by spreading information or bringing information out into the open will regularly reduce pluralistic ignorance, enable people to coordinate their actions and provide them with a reason to contribute to collective endeavours. Consequently, my epistemic obligations do not just concern *my own* beliefs, but those of others, too. Or, to put it more clearly, *my* epistemic obligations may concern *our* knowledge or beliefs. Sometimes, our obligations may require increasing interconnectedness of knowledge, sometimes symmetry of distribution or both (though there can be cases where asymmetry is conducive to a collective aim).

Having said that, since second-order beliefs are crucial for fostering shared efforts to address large-scale collective action problems, we may regularly need additional epistemic resources to reach a populace and instil public knowledge. Doing so will usually require some level of organization – typically governments will adopt mass information strategies, for instance, as part of public health campaigns, but non-governmental organizations (NGOs) and grassroots organizations such as Avaaz frequently resort to such measures, too. This takes me to the second claim defended here:

(2) Many of our epistemic obligations are probably best thought of as being collective in character.²³ The basic idea of collective moral responsibility or obligations is that these are held by groups of people either where (1) these form a constituted moral group agent²⁴ or where (2) sets of agents in so-called unstructured groups could collaborate towards a joint goal.²⁵ If epistemic obligations are – as Rosen argues – obligations to do certain things to improve our epistemic position and to take certain knowledge-enhancing action, then there should be no doubt that epistemic improvements are often more easily brought about and are often greater where they are the result of collaborative efforts. Epistemic enhancement, especially on a large scale, will regularly bear the features of joint necessity – that is, it will require collaboration of individual agents. In fact, if we think about the epistemic structure of our social world it is blindingly obvious that most of our greatest epistemic resources – libraries, archives, the entire bodies of scientific, medical as well as traditional and indigenous knowledge – are the result of continuing intergenerational epistemic collaboration.

In other words, the collective character of some of our epistemic obligations can be understood in at least two ways: (1) it seems reasonable to assume that so-called *group agents* (including states, corporations, governmental and NGOs) have obligations towards maintaining and expanding our joint epistemic resources. Such epistemic obligations of group agents are widely recognized and often formalized in the law. Arguably, though, epistemic obligations of group agents can extend beyond the legal requirements on them. Somewhat more controversially, (2) one might suggest that some of our epistemic obligations are held collectively by *sets of agents* where these sets do not themselves constitute a (novel) group agent. A number of scholars have argued for the existence of joint or collective *moral* obligations of this kind (Isaacs, 2011; Schwenkenbecher, 2019, 2021; Wringe, 2016).

It seems plausible to think that just like we sometimes have *moral* obligations to jointly perform or produce morally important actions or outcomes we can be required to jointly

produce epistemically important outcomes. In fact, doing the latter will often be the prerequisite for the former, as I have tried to explain above: In order to meet potential joint obligations we will regularly have to generate *joint epistemic goods*, such as ensuring that all agents in a collective action share the same true first- and higher order beliefs regarding their contributions.²⁶ Jointly generated epistemic goods will include goods that require several members of a set of agents to contribute, for instance, by providing other(s) with crucial pieces of information. In the aforementioned example of joint rescue the two potential helpers may need to share their assessment of the situation, of their own abilities, or of the best course of action. The aforementioned example also suggests that we should include epistemic goods that are *procedural* such as joint deliberation. And joint epistemic goods will include shared and common higher order beliefs where mutual reassurance is required in order to produce some further good, as in the examples given previously.

Note that in suggesting that epistemic obligations can be collective, I am not so much uncovering a novel type of obligation as I am making explicit something that we already – often intuitively – do. But since I do not have the space here to defend the notion of collective epistemic obligations in any detail, I will endorse only the following conditional claim: if (1) so-called group agents or (2) sets of agents can have collective *moral* responsibility or obligations, then we should accept that they can also have collective *epistemic* responsibility or obligations.

I will briefly point to what collective epistemic obligations might mean for individual agents who are members of those groups. Scholars of collective obligations agree that these entail contributory obligations for group members. When it comes to *group agents* as per (1), such as states and corporations, members' epistemic obligations are often an explicit and codified part of their role. For members of loose groups as in (2) – exemplified by our two rescuers above – these obligations would depend very much on the context. As I have indicated, we very regularly participate in the production of joint epistemic outcomes without giving that social practice much explicit thought. In the public domain, however, sharing our first-order or higher order beliefs with others in an effort to generate shared and public knowledge (or beliefs) may come less natural to us. The preceding discussion has hopefully succeeded in showing that it is important to speak up on issues that concern us,²⁷ not merely as a matter of acting with integrity and sincerity, but also as our contribution to generating accurate higher order beliefs in others and in our social group as a whole.²⁸ This will make it harder to be caught in what Bicchieri calls the 'belief-trap' – a scenario where we do not change harmful practices or do not change our behaviour so as to produce morally important collective goods (partly) because we are unaware that our privately held views are shared by many others in our group.

I started out with the question of when collective ignorance is blameless. I suggested that it is blameless when agents have met their epistemic obligations. I noted that epistemic obligations can be collective in nature and I briefly sketched what such obligations could entail for sets of agents in loose groups or for members of group agents. Beyond the explicit epistemic obligations that group agents hold by law or by virtue of their own internal constitution it will be difficult to establish in a general sense what such obligations entail and when they have been violated. The concrete collective epistemic obligations we hold as members of unstructured groups depend on our specific circumstances and are impossible to codify in the abstract. Since I do not commit to a substantive moral theory here, I cannot provide an answer to the question of the *extent* of our epistemic obligations (collective or individual) and how demanding they are. What I hope to have

shown, though, is that our epistemic obligations as members of such groups concern not just *our own* knowledge but that of others, too. We bear some responsibility for what others believe and we ought to strive to improve the epistemic status of our peer and social groups. Furthermore, such epistemic improvement will regularly require collective action.

Open Questions

In sum, when are we at fault for our collective ignorance and therefore blameworthy for harmful inaction resulting from our ignorance? The – somewhat preliminary – answer reached is this: ‘when we have violated our individual or collective epistemic obligations’. What this means exactly will need to be spelled out in more detail in future work. Unable to provide more than a sketch of such a theory here, I instead referred to existing work on collective *moral* obligations, which could serve as a starting point.

Let me briefly address two further open questions: that of choosing between individual and collective actions towards epistemic improvement and the question of blameworthiness for failing to meet our epistemic obligations. As for the first, discerning solutions to joint necessity problems *collectively* will often be better than trying to find solutions individually. However, initiating collaboration is also costly and at times risky. One of the issues we will come across in this context is the weighing up of how to best invest our limited epistemic resources. How much of an effort are we required to put into finding collective solutions to moral joint necessity problems? Salient, but imperfect solutions of which we have complete or at least sufficient knowledge will be competing against potentially superior solutions that require thorough investigation first. Furthermore, some of those epistemic resources may be individually available when others require a collective effort. When are we justified in acting on potentially limited, but readily available knowledge over investing more time in epistemically enhancing our collective resources? I cannot provide a satisfying answer to this question here, but will point to the parallel discussion concerning individual versus collective moral obligations in my recent book *Getting Our Act Together: A Theory of Collective Moral Obligations* (Schwenkenbecher, 2021: 63–68).²⁹ Suffice it to say that it is not obvious – as Doan seems to suggest – that we should *always* invest in improving our *joint* epistemic resources over addressing the moral problems we already have at least acceptable (though perhaps not perfect) individual solutions for. This would be a substantive claim that itself requires justification.

As for the second open question: if collective inaction resulting from group-based ignorance is blameworthy where collective epistemic obligations have been violated, this raises the question of ‘who is blameworthy?’ Again, this is an issue that I cannot devote any space to here. I have suggested elsewhere that, depending on the circumstances, blameworthiness for collective failings can be both individual (if an individual is solely responsible for the failing) or attach to the group or set of agents (Schwenkenbecher, 2021: 108–109).³⁰

And finally, we should keep in mind that even perfectly scrupulous moral deliberators can get it wrong sometimes and so there will always be cases where we cannot be blamed for our failure to discern perfectly knowable solutions.

Acknowledgements

The author feels lucky to have received extensive feedback on earlier versions of this article from a number of colleagues, including Suzy Killmister, Dave van Mill, Piero Moraro, Lachlan Umbers and Alan Tapper. She is

also very grateful to audiences at the following events, where this paper was presented: St Cross Ethics Seminar University of Oxford (2018); ACU International Conference on Moral Epistemology (2018); Freie Universität Berlin Workshop Politische Philosophie (2018); Kioloa Moral Philosophy Workshop (2019); Murdoch University Epistemology of Ignorance Workshop (2019); UWA Philosophy Workshop (2019); La Trobe University 4th Australasian Political Theory and Political Philosophy Conference (2020). She sincerely thanks two anonymous reviewers for this journal and the editor for excellent comments that greatly helped improve the paper. She gratefully acknowledges financial support from the Oxford Martin School at The University of Oxford where parts of this paper were written as part of an Oxford Martin Visiting Fellowship (2017–2018) and from the School of Arts at Murdoch University, which supported this research through the Research Study Leave Scheme (2017–2018).

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: Murdoch University's School of Arts (Research Study Leave Program) and Oxford Martin School at Oxford University (Oxford Martin Visiting Fellowship Program).

ORCID iD

Anne Schwenkenbecher  <https://orcid.org/0000-0003-2207-3043>

Notes

1. Besides Hormio's, there are two more articles on collective ignorance: (S. Miller, 2018; Ranalli and van Woudenberg, 2019). Rik Peels and Thirza Lagewaard (forthcoming) have a book chapter forthcoming on 'Group Ignorance. An account based on case studies of fundamentalist and white ignorance'. Major works on ignorance, even recent ones, largely bypass the issue of collective ignorance (DeNicola, 2017; Rescher, 2009).
2. Since many of the debates referenced here are couched in terms of 'knowledge', I will likewise use that term where the context requires it.
3. Seumas Miller (2018: 26) calls propositional ignorance 'non-doxastic ignorance'. Other types of non-doxastic ignorance include what Nadja El Kassar (2018) has called 'attitudinal ignorance'. I will restrict my exploration to doxastic ignorance while acknowledging that there is fertile ground for further investigation beyond those limits.
4. Factive ignorance includes cases where belief is consciously suspended as is the case where an agent knows that she does not know whether or not p is true (a known unknown). However, I will leave aside belief suspension cases here and focus on factive ignorance that obtains where agents have false beliefs.
5. At least we can safely assume that for 'ordinary' citizens. The greenhouse effect was, of course, already discovered in the nineteenth century by Svante Arrhenius, so many scientists should have been in the know.
6. I leave open what exactly counts as a large or small subset. This depends on the context. As such, widely and partially shared knowledge is somewhat imprecise notions (but deliberately so).
7. Fully shared ignorance is not defined simply as the failure of fully shared knowledge.
8. S. Miller (2018: 28) does not consider shared propositional ignorance as a type of collective ignorance, because there is no interconnection or interdependence between agents' absent doxastic 'states', as he puts it.
9. For the distinction between strict and wide joint necessity, see Schwenkenbecher (2019).
10. Some have criticized this definition as too demanding, but this criticism mainly targets definitions that do not restrict the number of iterations. One can avoid that by simply referring to 'n-level common knowledge' or 'common knowledge up to level n'. Other scholars, perhaps in order to avoid such criticism, have preferred to use the metaphor of knowledge being 'out in the open', such as Bratman (2014).
11. As a side note, some mobile phones now show a message's status as 'read' when it has been opened. Therewith, they produce higher order knowledge in both the sender and recipient. When the recipient

- reads a message, she knows that the sender knows that she read it because she knows that the sender will receive reading confirmation (or, to be exact, the sender can see whether *someone* has read the message).
12. Peels argues, *contra* LeMorvan, that propositional ignorance is just a subspecies of factive ignorance. I will not take sides in this debate, but simply note that Peels' point is compatible with my argument here.
 13. Floyd Allport is usually credited with coining this term (Katz and Allport, 1931). The literature on this issue is mainly in psychology (see, for example, Geiger and Swim, 2016; Latané and Darley, 1970; D. T. Miller and McFarland, 1987; Prentice and Miller, 1996). So far, there are few philosophers to discuss it (see, for example, Bicchieri, 2017; Rendsvig, 2014).
 14. However, Bicchieri and Mercier do not think changing second-order beliefs is *sufficient* for changing behaviour (2014: 44–45).
 15. They do, however, include a caveat: correcting pluralistic ignorance is only effective where the ignorance concerns the audience one is addressing (Geiger and Swim, 2016: 88). This seems to reflect Bicchieri's (2017) point that social norms function within specific reference networks.
 16. Of course, people can shift their behaviour in ways that they may not actively choose. 'Nudging' – or choice architecture – in particular can help shift behaviour while bypassing conscious decision-making (Thaler and Sunstein, 2008). Take the example of making organ donation an opt-out rather than an opt-in decision with a view to increasing the pool of organ donors. However, in the absence of such choice manipulations, many collective action problems will only be addressed if people *decide* to perform certain actions or to change their behaviour.
 17. Like Held, Felix Pinkert (2014) argued that collective obligations depend on there being a *salient* solution to a collective moral action problem such that it is immediately obvious what each individual must do in order to contribute. Pinkert also thinks we can have (mediate) collective obligations where we first need to establish joint ability by working out some collective action plan. According to Tracy Isaacs (2011: 152), '[o]nly when the course of action presenting itself is *clear* to the *reasonable* person is it accurate to think in terms of the collective obligations of putative groups' (my emphasis). At a different point she speaks of solutions coming 'into focus' (Isaacs, 2011: 140).
 18. This is true of Held, Isaacs and Pinkert, but also of May (1992) and Petersson (2008).
 19. However, the authors critiqued by Doan by no means claim to be covering all or even the most important types of inaction (and its link to ignorance). Furthermore, Held (1970) thinks that a set of agents may still bear (negative) responsibility for not turning themselves into the kind of collective that is capable of making collective decisions.
 20. That is, I deliberately bypass the discussions on group rationality and collective decision-making, both because I do not have enough space and because they are not relevant to my argument. See, for instance, Kopec (2019).
 21. Rosen (2004: 307) ultimately argues that '*the only possible locus of original responsibility is an akratic act*', because that is what it would take for an agent to act badly without ignorance. That is, she knows it is wrong and that all-things-considered she should not do it, but she does it regardless. On Rosen's (2004: 308–310) view, this leads to a sceptical argument about our ability to ascribe responsibility for any particular act of wrongdoing: we can never know for certain whether an act resulted from genuine akrasia rather than just ordinary weakness of will. This sceptical upshot need not concern us here, though.
 22. There are countless publications on this topic. For an overview, see Bazargan-Forward and Tollefsen (2020), Hess et al. (2018) and May and Hoffman (1991).
 23. The collective epistemic obligations I have in mind are different from Sandy Goldberg's (2018: 186) *social epistemic responsibilities*, which concern the extent to which one meets the epistemic expectations of others in the formation of one's beliefs. My position, in contrast, focuses on our obligations to generate higher order beliefs in others – often simply by confirming our beliefs to them – in order to prevent or fix pluralistic ignorance, or, more simply to enable and facilitate collective action and coordination.
 24. On moral group agents, see List and Pettit (2011) or Tollefsen (2015).
 25. For an overview of existing accounts of collective obligations, see Schwenkenbecher (2021, chapter 6).
 26. Margaret Gilbert's (2004) 'collective belief' would constitute a joint epistemic good. A further joint epistemic good is the collective prevention of 'epistemic harms' (Fleisher and Šešelja, 2020).
 27. For further discussions of our obligations to speak up see Jennifer Lackey on 'duties to object' and Katherine Furman on 'suppressed disagreement' (Furman, 2018; Lackey, 2020).
 28. See also Boyd Millar (2019) on our shared responsibility to counter misinformation.
 29. Another relevant debate is that on trade-offs between pattern-based reasons or we-reasoning and individual-based reasoning in Christopher Woodard's (2011, 2017) work.
 30. See also Millar (2019).

References

- Bazargan-Forward S and Tollefsen D (2020) *Handbook of Collective Responsibility*. London: Routledge.
- Bicchieri C (2017) *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. New York: Oxford University Press.
- Bicchieri C and Mercier H (2014) Norms and Beliefs: How Change Occurs. In: Xenitidou M and Edmonds B (eds) *The Complexity of Social Norms*. Cham: Springer International Publishing, pp.37–54.
- Bratman ME (2014) *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press.
- DeNicola DR (2017) *Understanding Ignorance the Surprising Impact of What We Don't Know*. Cambridge, MA: The MIT Press.
- Doan MD (2016) Responsibility for Collective Inaction and the Knowledge Condition. *Social Epistemology* 30 (5–6): 532–554.
- El Kassir N (2018) What Ignorance Really Is. Examining the Foundations of Epistemology of Ignorance. *Social Epistemology* 32 (5): 300–310.
- Fleisher W and Šešelja D (in press) Collective Epistemic Responsibility: A Preventionist Account. Epub ahead of print 14 March 2020. Available at <http://philsci-archive.pitt.edu/17003/>
- Funk S, Knapp JK, Lebo E, et al. (2019) Combining Serological and Contact Data to Derive Target Immunity Levels for Achieving and Maintaining Measles Elimination. *BMC Medicine* 17 (1): 180.
- Furman K (2018) Moral Responsibility, Culpable Ignorance and Suppressed Disagreement. *Social Epistemology* 32 (5): 287–299.
- Geiger N and Swim JK (2016) Climate of Silence: Pluralistic Ignorance as a Barrier to Climate Change Discussion. *Journal of Environmental Psychology* 47: 79–90.
- Gilbert M (2004) Collective Epistemology. *Episteme* 1 (2): 95–107.
- Godman M (2013) Why We Do Things Together: The Social Motivation for Joint Action. *Philosophical Psychology* 26 (4): 588–603.
- Goldberg S (2018) *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford: Oxford University Press.
- Held V (1970) Can a Random Collection of Individuals Be Morally Responsible? *Journal of Philosophy* 67 (14): 471–481.
- Hess K, Igeneski V and Isaacs TL (eds) (2018) *Collectivity: Ontology, Ethics, and Social Justice*. London: Rowman & Littlefield.
- Hormio S (2018) Culpable Ignorance in a Collective Setting. In: Kuorikoski J and Toppinen T (eds) *Action, Value and Metaphysics – Acta Philosophica Fennica*, vol. 94. Helsinki: Philosophical Society of Finland, pp.7–34.
- Isaacs TL (2011) *Moral Responsibility in Collective Contexts*. Oxford: Oxford University Press.
- Katz D and Allport FH (1931) *Students' Attitudes: A Report of the Syracuse University Reaction Study*. Syracuse, NY: Craftsman Press.
- Kopec M (2019) Unifying Group Rationality. *Ergo: An Open Access Journal of Philosophy* 6: 517–544.
- Lackey J (2020) The Duty to Object. *Philosophy and Phenomenological Research* 101 (1): 35–60.
- Latané B and Darley JM (1970) *The Unresponsive Bystander: Why Doesn't He Help?* New York: Appleton-Century Crofts.
- Le Morvan P (2011) On Ignorance: A Reply to Peels. *Philosophia* 39 (2): 335–344.
- List C and Pettit P (2011) *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford; New York: Oxford University Press.
- Louis WR, Taylor DM and Neil T (2004) Cost-Benefit Analyses For Your Group And Yourself: The Rationality Of Decision-Making In Conflict. *International Journal of Conflict Management* 15 (2): 110–143.
- May L (1992) Collective Inaction and Responsibility. In: May L (ed.) *Sharing Responsibility*. Chicago, IL: Chicago University Press, pp.105–124.
- May L and Hoffman S (1991) *Collective Responsibility: Five Decades of Debate in Theoretical and Applied Ethics*. Savage, MD: Rowman & Littlefield.
- Millar B (2019) Shared Epistemic Responsibility. *Episteme*. Epub ahead of print 28 August. DOI: 10.1017/epi.2019.21.
- Miller DJ (2017) Reasonable Foreseeability and Blameless Ignorance. *Philosophical Studies* 174 (6): 1561–1581.
- Miller DT and McFarland C (1987) Pluralistic Ignorance: When Similarity Is Interpreted as Dissimilarity. *Journal of Personality and Social Psychology* 53 (2): 298–305.

- Miller S (2018) Collective Knowledge and Collective Ignorance. In: *Dual Use Science and Technology, Ethics and Weapons of Mass Destruction*. Cham: Springer, pp. 21–37.
- Peels R (2012) The New View on Ignorance Undefeated. *Philosophia* 40 (4): 741–750.
- Peels R and Lagewaard T (forthcoming) Group Ignorance. An Account Based on Case Studies of Fundamentalist and White Ignorance. In: Lackey J and McGlynn A (eds) *Oxford Handbook of Social Epistemology*. Oxford: Oxford University Press.
- Petersson B (2008) Collective Omissions and Responsibility. *Philosophical Papers* 37 (2): 243–261.
- Pinkert F (2014) What We Together Can (Be Required To) Do. *Midwest Studies in Philosophy* 38 (1): 187–202.
- Prentice DA and Miller DT (1996) Pluralistic Ignorance and the Perpetuation of Social Norms by Unwitting Actors. *Advances in Experimental Social Psychology* 28: 161–209.
- Ranalli C and van Woudenberg R (2019) Collective Ignorance: An Information Theoretic Account. *Synthese*. Epub ahead of print 13 September. DOI: 10.1007/s11229-019-02367-7.
- Rendsvig RK (2014) Pluralistic Ignorance in the Bystander Effect: Informational Dynamics of Unresponsive Witnesses in Situations Calling for Intervention. *Synthese* 191 (11): 2471–2498.
- Rescher N (2009) *Ignorance: (On the Wider Implications of Deficient Knowledge)*. Pittsburgh, PA: University of Pittsburgh Press.
- Rosen G (2003) IV – Culpability and Ignorance. *Proceedings of the Aristotelian Society (Hardback)* 103 (1): 61–84.
- Rosen G (2004) Skepticism about Moral Responsibility. *Philosophical Perspectives* 18 (1): 295–313.
- Roy O and Schwenkenbecher A (2019) Shared Intentions, Loose Groups, and Pooled Knowledge. *Synthese*. Epub ahead of print 8 August. DOI: 10.1007/s11229-019-02355-x.
- Schwenkenbecher A (2019) Collective Moral Obligations: ‘we-Reasoning’ and the Perspective of the Deliberating Agent. *The Monist* 102 (2): 151–171.
- Schwenkenbecher A (2021) *Getting Our Act Together: A Theory of Collective Moral Obligations*. London; New York: Routledge.
- Singer P (1972) Famine, Affluence, and Morality. *Philosophy and Public Affairs* 1 (3): 229–243.
- Thaler RH and Sunstein CR (2008) *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.
- Thomas EF, McGarty C and Mavor KI (2009) Aligning Identities, Emotions, and Beliefs to Create Commitment to Sustainable Social and Political Action. *Personality and Social Psychology Review* 13 (3): 194–218.
- Thomas EF and Louis WR (2013) Doing Democracy: The Social Psychological Mobilization and Consequences of Collective Action. *Social Issues and Policy Review* 7 (1): 173–200.
- Tollefsen D (2015) *Groups as Agents*. Cambridge: Polity Press.
- Woodard C (2011) Rationality and the Unit of Action. *Review of Philosophy and Psychology* 2 (2): 261–277.
- Woodard C (2017) Three Conceptions of Group-Based Reasons. *Journal of Social Ontology* 3 (1): 107–127.
- Wringe B (2016) Collective Obligations: Their Existence, Their Explanatory Power, and Their Supervenience on the Obligations of Individuals. *European Journal of Philosophy* 24 (2): 472–497.

Author Biography

Anne Schwenkenbecher is Senior Lecturer in Philosophy at Murdoch University in Western Australia. Her research areas include collective agency and responsibility, social epistemology, environmental ethics and political violence. She is the author of *Terrorism: A Philosophical Enquiry* (Palgrave Macmillan, 2012) and *Getting Our Act Together: A Theory of Collective Moral Obligations* (Routledge, 2021).