


## ARTICLE

# The Collective Responsibilities of Science: Toward a Normative Framework

Vincenzo Politi 

Institute for the History of Science, Universitat Autònoma de Barcelona, Bellaterra, Barcelona, Spain  
Email: [vincenzo.politi@uab.cat](mailto:vincenzo.politi@uab.cat)

(Received 12 September 2023; revised 04 June 2024; accepted 19 June 2024; first published online 16 September 2024)

## Abstract

Scientists have the epistemic responsibility of producing knowledge. They also have the social responsibility of aligning their research with the needs and values of various societal stakeholders. Individual scientists may be left with no guidance on how to prioritize and carry these different responsibilities. As I will argue, however, the responsibilities of science can be harmonized at the collective level. Drawing from debates in moral philosophy, I will propose a theory of the collective responsibilities of science that accounts for the internal diversity of research groups and for their different responsibilities.

## 1. Introduction

Scientists have *epistemic responsibilities*, encompassing the duties to make, justify, and transmit true claims about their domain. They also have *social responsibilities*, encompassing the duties to align their research toward socially desirable ends and democratically held values. Emerging science policies, such as the Ethical, Legal, and Social Implications framework (ELSI) or Responsible Research and Innovation (RRI), seek to implement social responsibility in scientific research groups. Among their proposed strategies, there are public engagement, the encouragement of interdisciplinary interactions between STEM researchers and scholars from the humanities and the social sciences, and fostering researchers' anticipatory and reflecting attitude about the potential social and cultural impacts of their work.

Epistemic and social responsibilities may pose some difficulties to scientists, who are left without guidance for prioritizing different duties (Resnik and Elliott 2016). Some activities regarded as promoting the alignment of science with society, such as public engagement, may be too time consuming, thus becoming a potential obstacle to scientists' other duties. Many scientists may not be used to answering challenging questions from concerned nonexperts and they may even feel uneasy doing so. Depending on how they engage with the public, some may even put their reputation at risk (Pielke 2007). Such a risk is not counterbalanced by an appropriate reward system for performing this kind of duty.

© The Author(s), 2024. Published by Cambridge University Press on behalf of the Philosophy of Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Scientists, however, are not lone agents. Scientific knowledge is produced by research groups and scientific communities. Science, in other words, possesses a collective dimension. I argue, therefore, that the primary bearers of the responsibilities of science are not the individual scientists, but the scientific groups. I argue, furthermore, that the problems that carrying out different responsibilities may cause to individual scientists could be solved at the collective level.

In section 2, I introduce some concepts borrowed from moral philosophy to define the social responsibility of science in terms of *collective remedial responsibility*. In section 3, I discuss how philosophy of science often fails to properly discuss the collective dimension of the responsibilities of science. On the one hand, individualistic theories do not account for how different individual scientists occupy different positions within a research group, and for how collective responsibility is achieved through the internal diversity of groups. On the other hand, social approaches remain unclear about the structure of well-designed responsible research groups. In section 4, I assess two of the few accounts of the collective responsibility of science—the *duty to join forces* account (Fleisher and Šešelja 2022), and the *special roles* account (Rolin 2017). Even though these accounts are only about the collective *epistemic* responsibility of science, and not about the *relation* between collective epistemic and collective social responsibility, it is still worth discussing their merits and limitations, as well as the way in which they could be amended to account for the different responsibilities of science. As it turns out, neither account is satisfying under this respect. The duty to join forces account risks imposing too strenuous demand on individual scientists and leaves it unclear exactly how relevant individuals should join forces, while the special roles account threatens to collapse collective responsibility into individual responsibility. In section 5, I advance my own proposal, which builds on the two previous accounts and overcomes some of their problems. In my framework, to carry out all the different responsibilities, a research group ought to distribute *both* special duties *and* collective duties. In section 6, I discuss how my normative framework could work for scientific groups of different levels and sizes.

## 2. The social responsibility of science as collective responsibility

“Social responsibility of science” is an umbrella expression that subsumes different meanings and can be interpreted and operationalized in different ways (Timmermans and Blok 2021; de Melo-Martín and Intemann 2023). To understand why it is different from epistemic responsibility, however, it is necessary to define it. For this reason, in this section I refer to some of the concepts developed by moral philosophers to bring some clarity on the minimal commonalities shared by the different versions of the concept of social responsibility of science.

First of all, what is responsibility? Moral philosophers distinguish between *backward-looking* and *forward-looking responsibility*, as well as between *individual* and *collective responsibility* (see Smiley 2022). Backward-looking responsibility amounts to either accountability or blameworthiness, and it is at the root of the legislative views on the post hoc evaluation of wrongdoings and eventual assignment of punishment. Forward-looking responsibility is about the duties to produce a morally desirable outcome. Forward-looking responsibility can be either *preventionist* (related to the duty of preventing future harms) or *remedial* (related to the duty of bringing forth

positive changes). Backward- and forward-looking responsibility are complementary; for example, a post hoc blame or punishment is assigned to someone who has failed to carry out their duties. Therefore, many theories of responsibility account for both.

The concepts of individual and collective responsibility are about the bearers of responsibility. Very broadly speaking, to be attributed responsibility, an agent must meet the *epistemic condition* (she ought to have the right kind of knowledge and awareness), the *freedom (or control) condition* (she ought to be able to act free from external pressures), and the *intentionality condition* (she ought to intend to act in a certain way, meaning that responsibility is not attributed to someone who, e.g., ends up doing the right thing unintentionally, or even out of sheer luck; for an overview on these concepts, see Talbert 2023; Rudy-Hiller 2022; Mele and Sverdlik 1996). Some philosophers argue that these conditions are not met exclusively by individual people.

For the supporters of the concept of collective responsibility, the knowledge possessed by a group exceeds the sum of the bits of knowledge possessed by its members. They argue, moreover, that groups may perform actions that individuals alone would be incapable of doing. Finally, they also argue that members of a group performing a collective action are all contributing to reach a collective objective, which, however, may be different from their own individual ends: In this sense, groups can also meet the intentionality condition. If one accepts that groups possess knowledge, can act, and have intentions, and also that group-knowledge, group-actions, and group-intentions are irreducible to the simple aggregation of what their individual members know, of how they act, and of what they intend, then one may also accept that responsibility can sometimes be attributed to collectives, such as groups of people, corporations, professional bodies, or nations.

Due to the irreducibility of collective agents, supporters of collective responsibility also theorise about a series of asymmetries between groups and individuals. For example, some argue that, in some cases and under certain conditions, responsible individuals may end up forming a group that behaves irresponsibly, and vice versa (Coop 2007).

Not every philosopher accepts the existence of collective agents. Individualistic arguments against collective responsibility stress that groups do not have a mind of their own and, therefore, cannot meet the intentionality condition for collective agency (Lewis 1948; Watkins 1957; Goldman 1970; Sverdlik 1987; Corlett 2001; Narveson 2002).

Nonindividualistic approaches, however, do not support controversial metaphysical views about the existence of a rather mysterious “collective mind.” Rather, they focus on the structure a group ought to possess to be a candidate for the attribution of collective responsibility. For Gilbert (1990, 2000, 2013) a *plural subject* is established when different individuals are unified by a “joint commitment” (i.e., when all the individuals share the same aim, they know that they all share the same aim, and act accordingly). For Bratman (1992, 1993, 2013), however, it is possible to imagine a group of individuals who meet the conditions of joint commitments but who do not communicate or cooperate with one another. Such a group could hardly be considered as a collective agent. In Bratman’s view, therefore, collective intentionality can be attributed only to *collective subjects*, whose members are engaged in a “shared cooperative activities.”

Finally, philosophers such as French (1984), Rovane (1997), List and Pettit (2011), and Collins (2019) argue, each in their own ways, that collective responsibility can be attributed only to those groups operating under a group-level decision system, the output of which cannot be reduced to the aggregation of individual preferences. Such groups are internally organized and may assign different duties to different individual members. Another characteristic of these organized groups is not only the ability of their members to coordinate their actions to reach a collective goal, but also their ability to set *new* goals and even to *change* them.

Discussing in detail all the arguments for and against collective responsibility, as well as all the different views about joint intentions and shared duties within groups, would go far beyond the limits of this article. For the present purposes, it is enough to stress how the debates in moral philosophy offer a wealth of resources to reflect upon some issues in philosophy of science, such as the definition of the responsibilities of science and the structure of responsible collectives.

Since the seminal works of John Dewey, Ludwig Fleck, Thomas Kuhn, and many others, philosophers of science regard the scientific community, instead of the (ideally rational) individual, as the knowledge-producing agent. Several philosophers argue that not only does the scientific community produce knowledge but it also possesses it (see, e.g., Bird 2010, 2022; Miller 2015). Others maintain that science as an institutionalized collective has its own aims, which may differ from the aims of its individual members (Bird 2019, 2022). It has also been argued that the public trust on science is not directed toward individual researchers, but to collective bodies, which are therefore the bearers of responsibility (Wilholt 2016). From these arguments, it appears like scientific communities can be considered as collective agents. Accordingly, philosophers have developed some accounts of the epistemic responsibility of science as a kind of collective responsibility. Similarly, it is possible to regard the social responsibility of science as a kind of collective responsibility too.

To be more specific, it must be noticed that extreme cases of scientists' gross misconduct are prosecutable. In other words, there exist juridical mechanisms to hold accountable, blame, and eventually punish scientists in an after-the-fact fashion. The social responsibility of science promoted by emerging science policies, by contrast, is about scientists' duties to work toward socially desirable ends. Therefore, it is a form of *forward-looking responsibility*.

Furthermore, professional codes of research integrity as well as ethics committees are already in place, at the national and international level, and provide the rules that scientists ought to follow to prevent harm to human and animal subjects, as well as to society at large, and to regulate so-called dual-use. While these regulations may not be enough to prevent all kinds of harm (e.g., unregulated "gray areas" may still exist, especially when it comes to innovative technologies), the social responsibility of science seems to be about something more than harm prevention. Such a concept has become central in the science policy discourse because of the recognition that the potential implications of scientific research are not always clear-cut and somehow quantifiable harms, and may involve cultural and social transformations, or "soft impacts" (van der Burg 2009). Philosophers have also noticed that science and technology may have not only unintended practical consequences but they may also promote or transform some values unintendedly (Ratti and Russo 2024).

In short, it appears that the social responsibility of science currently discussed in science policy and philosophy of science can be defined as a form of *collective remedial responsibility*. This kind of responsibility adds to epistemic responsibility.

Having defined the social responsibility of science as a form of collective remedial responsibility, it remains to investigate how a scientific research group ought to be structured to be the bearer of both epistemic and social responsibility.

### 3. Responsible scientists or responsible communities?

Some philosophers of science seem to fail to address the collective character of the responsibilities of science. For example, some supporters of the so-called argument from inductive risk, such as Rudner (1953) and Douglas (2009), frame the responsibilities of science in individualistic terms. The argument states that because their errors may have harmful consequences and they have the same responsibilities as everybody else, scientists ought to make value judgments to mitigate the risk of error in the face of the available evidence.

In Rudner's formulation, "[I]n accepting a hypothesis the scientist must make the decision that the evidence is sufficiently strong or that the probability is sufficiently high to warrant the acceptance of the hypothesis" (Rudner 1953, 2). More synthetically, and to recall the title of his widely read article, for Rudner "the scientist *qua* scientist" must make value judgements. This reference to the individual scientist as the decision-making agent and, therefore, as the bearer of responsibility could be read as a way of conceiving research groups and scientific communities as "individuals writ large."

While it is not clear whether Rudner held such an individualistic outlook, Douglas clearly stated that, although she appreciates the philosophical analysis of the social dimensions of science, "we also need clear norms for *individual* reasoning in science" (Douglas 2009, 19; original emphasis). More recently, she has also claimed that although "[a]ll scientists must attend to the question over whether they *should*, not just whether they *can*," in her view, "[u]ltimately, freedom and responsibility together must rest with the individual scientist" (Douglas 2021, 83).

A significant problem for this individualistic outlook, though, is that it seems to take for granted that it is possible to isolate an individual's reasoning and responsibilities from the group of which they are part. Some have argued, furthermore, that the individualistic approach revolves around an ideal responsible scientist, a fictional individual who does not help much in capturing the complexity of collective dynamics. For example, Lefevere and Schliesser (2014) discuss scenarios in which responsible individual scientists end up making bad or harmful decisions by following their community standards. Cases like this urge us to look at the collective dimension of the formulation of decision-making community standards and methods. Politi and Grinbaum (2020) borrow some insights from social epistemology and claim that "ethical labor," like its cognitive counterpart, is distributed among different individual members of the scientific community.

Philosophers like Longino (1990, 2002), Kitcher (2001, 2011), and several others have developed normative models for responsible scientific communities. They suggest that the community of inquirers should be open to receive and uptake public scrutiny (as in the case of Longino's value management ideal), or that it should be

democratized through the active engagement in an ideal rational conversation with some instructed members of the public (as in the case of Kitcher's Well-Ordered Science). Contrary to those following an individualistic approach, for these philosophers the responsible agent in science is the scientific community, rather than its individual members.

More recently, Janet Kourany (2021) has explicitly stated that the issue of the social responsibility of science requires a "shift from an individual perspective to a social perspective" (106). Kourany makes an analogy between workers' unions and scientific communities: in the same way in which single workers would not have the power to change their conditions on their own but would be able to exert some form of control through a collective union, so the scientific community as a whole, if driven by the right values, may bring positive social change in a way that individual scientists alone would not be able to achieve. Indeed, "[I]f the scientific community as a community were to refuse to pursue a certain kind of research . . . , or were to insist on a different kind of research . . . , such a stand could effectively control the research that is done and its aftermath in a socially responsible way" (ibid.).

The problem with these proposals is that they do not discuss how a well-designed responsible scientific group ought to be structured. One thing is to say that the scientific community ought to engage with democratic institutions, the public, and so on; another thing is to specify how it ought to be organized in such a way that its engagement with democratic institutions or the public is facilitated. For example, *who* or *how many*, within the scientific community, ought to carry the duties of engaging with other stakeholders or responding to public criticisms? Following on Kourany's analogy, workers' unions are institutionalized collectives organized in such a way that few selected individuals may engage with other institutions and represent the interests of the whole group. What many accounts of a responsible scientific community miss, however, is precisely a reflection on its internal structure and organization.

Philosophers of science, in short, are divided on the issue of the responsible agent of science. Some of them attribute such a responsibility to individual scientists, others to groups. The two different perspectives lead to incompatible normative frameworks: For instance, what is considered to be acceptable for a group would be deemed as unacceptable for an individual (Schroeder 2022). However, because they do not explicitly acknowledge and discuss the differences in their approaches, philosophers debating about the responsibilities of science may end up talking at cross-purposes.

Moreover, both individual and collective accounts of scientific responsibility fail to adequately consider the differentiated roles that scientists may have in their communities. These different roles may correspond to different responsibilities. In other words, present philosophical discussions about the responsibilities of science, whether individualistic or collective, do not account for the fact that not every individual scientist may carry all the duties at the same time and in the right way, nor do they account for issues concerning the internal structure of a responsible scientific community.

Without a reflection on these issues, both the discussions about the responsible scientist and the responsible scientific community are incomplete as a basis for the development of a normative theory of the responsibilities of science. Such a theory could be modeled after some of the theories about collective responsibility developed by moral philosophers.

#### 4. The collective responsibilities of science: Two theories

Not many philosophers of science have drawn from the moral philosophy literature on collective responsibility to develop normative theories on the structure of responsible scientific groups. Two notable and recent exceptions are Fleisher and Šešelja (2022; from now on, FS) and Rolin (2017).

It must be specified that both FS's and Rolin's theories are about the collective *epistemic* responsibility of science, whereas the aim of this article is to account for how epistemic *and* social responsibilities may be harmonized within groups. Nevertheless, FS's and Rolin's proposals still deserve to be analyzed attentively, especially to assess whether and how they account for science as a responsible collective agent.

##### 4.1 Individual duties and mobilized groups

While many philosophers interested in collective epistemic responsibility analyze the course of actions leading to the formation of group beliefs, FS (2022) are concerned with how to prevent collective epistemic harms, that is, epistemic harms that affect a collectivity and could be prevented through collective action.

To explain their view, FS use of a rescue scenario:

*Rescue #1:* A group of twelve children are swimming in the ocean. Three of the children 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. (FS 2022, 8)

Following Hindriks (2019), the responsible course of action for random, unorganized groups, such as that described Rescue #1, consists of a two-step process and comprises both individual and collective duties. In the first step, individuals ought to mobilize the others, who in turn ought to join forces. In the second step, the mobilized group has the collective duty to prevent harm. The second duty is conditional on the fulfilment of the first: The random collective ought first to become a “mobilized group” to save all the children. Collective harm is successfully prevented by completing the two-step process. It may, however, occur for the failure of completing either the first or the second step.

FS apply Hindriks's framework for preventing collective moral harms to discuss the prevention of collective epistemic harms in science. They recognize that, while Hindriks focuses on harms afflicting random crowds, “scientific institutions act as organised groups that aim at both promoting epistemic goals and preventing epistemic harms” (FS 2022, 2). Nevertheless, they also contend that scientific communities may be unprepared to prevent some specific kinds of epistemic harm. This would be the case, for instance, of implicit methodological biases held by the majority of scientists, or when a potentially valid hypothesis risks to be prematurely abandoned not because the available evidence conclusively speaks against it, but for other contingent reasons (as it happened in the case discussed by Šešelja and Straßer 2014). When such circumstances arise, scientific communities too could be regarded as unorganized groups. In these cases, individual scientists have the individual duty of mobilizing their colleagues, who in turn have the individual duty of joining forces. After a sufficient number of scientists



have joined forces, the mobilized group will have the collective responsibility of preventing the collective epistemic harm.

FS's account is forward-looking in an explicit preventionist way: It is concerned with the avoidance of future harms. As specified in the introduction, scientists have both epistemic and social responsibilities; furthermore, as explained in section 2, the social responsibility of science is *remedial*, rather than preventionist. Nevertheless, it is worth considering whether FS's theory could be extended to account for collective remedial responsibility. One may argue, in fact, that the collective action driven by a sense of remedial responsibility is all the more successful when individuals manage to convince others to join forces in creating mobilized groups or social movements. *Prima facie*, therefore, there seems to be no particular obstacle to the extension of FS's preventionist framework to a theory of social remedial responsibility.

There is, however, an important disanalogy between *Rescue #1* and the scenarios FS have in mind. *Rescue #1* involves a small group of people who all have a clear understanding of the harm they must prevent. The cases FS seek to explain involve wide scientific communities. The bigger the size of the group to be mobilized, the higher the chances of incurring in the so-called *problem of insignificant hands*. When a morally desirable outcome can only be produced by many people, their individual actions, by themselves, will not produce it. In such cases, members of a very large group may feel that their actions are morally insignificant and, therefore, may avoid contributing. The problem of insignificant hands involves the first step of Hindriks's account in case it is applied to very large groups. If every individual thought that their joining forces will have no real impact in the formation of a mobilized group, then to begin with such a mobilized group will not be formed.

Moral philosophers have proposed a spectrum of solutions to this problem. On the one side of the spectrum, Sinnott-Armstrong (2005) argues that *nobody* has the duty to contribute to a collective preventionist action because nobody's action on its own can prevent collective harm, which means that nobody has that responsibility. This view is clearly problematic because it makes it difficult to understand how any collective harm could ever be prevented. On the other side of the spectrum, Kagan (2011) claims that *everybody* ought to contribute to a preventionist action, even when there is no chance that the desirable outcome will be obtained. This view makes it difficult to understand why anybody would carry out a duty knowing that no good outcome will come out of it. Hindriks (2022) has recently proposed his own solution to the problem of insignificant hands, based on the idea that individuals are required to contribute to a collective action, and indeed join forces, when the "prospect of success" of their actions is good enough.

Although they do not tackle it, FS's account also faces the problem of insignificant hands. Let's assume that a scientific community faces the risk of prematurely abandoning a valid hypothesis. Within such a community, it is plausible to suppose that there are individuals with different degrees of knowledge about X (the hypothesis that risks being unjustly abandoned), Y (likelihood of X being abandoned), and W (the harm caused by the abandonment of X). It is hard to imagine that everybody has the same amount of knowledge about X, Y, and W, especially if one is more familiar with, for example, hypothesis J.

Different scientists, furthermore, occupy different positions of power, with the more knowledgeable and more powerful presumably carrying more responsibilities.



Who would then be the most appropriate person to attempt to mobilize others? The postdoc who needs to publish the article by the end of the month or the world-known professor?

As already discussed, some moral philosophers would say that none of the scientists has the duty to join forces. In this way it will be hard to see how a mobilized group for preventing a collective epistemic harm could ever be formed. Others would maintain that the duty to join forces and mobilize others is everyone's responsibility. However, it is difficult to understand why junior scientists with little knowledge about X, Y, and W and little power ought to give insignificant contributions in mobilizing others. If we follow Hindriks (2022), the "prospect of success" is higher when the mobilizing action is performed by scientists with more knowledge about X, Y, and W, and with higher power (i.e., established scientists with an international reputation). Scientists with not enough knowledge about X, Y, and W, and with not enough power would therefore be dispensed with the duty of joining forces. This means that, in very big groups, such as scientific communities, the first step of mobilization involves only a selected number of individuals who ought to act in an organized way to decide who and how many ought to be mobilized.

In the case of the prevention of epistemic harms, like those discussed by FS, such a degree of group organization may not be too difficult to achieve after all. Not only are scientists trained to be epistemically responsible, but they are able to recognize who, among them, are the most knowledgeable and recognized personalities, that is those who ought to carry special epistemic duties, such as the duty of joining forces in forming a mobilized group to prevent collective epistemic harms. The problem, however, becomes more complicated in the case of social responsibility. Spotting the right individuals to carry group duties related to social responsibility may be more difficult. Moreover, the "right" people to carry epistemic responsibilities may not be the "right" people to carry social responsibilities.

## 4.2 Special duties

Rolin (2017) shows how a scientific community can be collectively responsible even though not all its individual members are required to carry the same duties. She does so by developing a theory of the distribution of special responsibilities within the scientific community. As she explains, while *general responsibilities* are universal, *special responsibilities* are carried by individuals who are assigned a particular role, thanks to the special relations they have developed with other individuals within their community, or with other communities.

To explain the difference between general and special responsibilities, Rolin relies on yet another rescue scenario.

*Rescue #2:* People who happen to be on a beach on a particular day all have a general moral duty to rescue a person who is in danger of drowning. For many reasons, not all persons on the beach are actually capable of carrying out this duty. If they all attempted to meet the duty to rescue the person in the water, the outcome would be a chaotic situation where more than one person is in danger of being drowned. A better way to implement a general moral duty is to assign it to a person who is well suited for the task of rescuing people in water:

the lifeguard. When a general moral duty is assigned to a particular person standing in a particular relation, it gives rise to a special moral duty that the person has. Thus, the lifeguard has a special moral duty to help people in life-threatening situations on the beach. (Rolin 2017, 477)

It is worth noticing that *Rescue #2* nicely exemplifies the asymmetry between individual and collective responsibility mentioned in section 2. Rescuing a person in danger is a general responsibility: Everybody ought to do so. Nevertheless, when someone is about to drown, not everybody should go to the rescue. Many uncoordinated and nonprofessional swimmers could get in each other's way and, as a result, they may end up not saving the person in danger, or even drawing themselves. In this case, the aggregated action of well-intentioned individuals with shared intention could result in harm. To avoid undesirable consequences, the professional lifeguard is designated to carry out the special responsibility of diving into the sea and saving people in danger. In a similar way, Rolin argues that only some individual scientists will have to carry special responsibilities for the whole scientific community.

The strength of Rolin's proposal can be appreciated when considering cases of scientific cooperation in which scientists do not act for their own immediate individual interests. As Rolin points out, some scientific debates require the contribution of experts from other specialties to be resolved. When they take part in debates outside their special field, scientists do not make their own community advance toward its epistemic objectives, nor are they serving their own individual interests. At the end of the collaboration, they will likely get just an acknowledgment in someone else's article, or they will be offered the co-authorship in articles that will have little or no impact in their own field, and for which they will receive little or no reward from their own community. Scientists have the general epistemic responsibility of contributing to the progress of knowledge but expecting that every single scientist of a scientific community contributes to debates in fields outside their own would be far too demanding. Such a scenario could even be harmful both for the individual scientists, who will sacrifice their personal career, and for their own scientific community, which will fail to make progress because too many of its members are busy contributing to foreign debates. Only individuals who have developed a "special relation" with the members of other groups will have to carry the special duty of contributing to their debates.

While Rolin's theory does not rely on the questionable idea that every single scientist ought to perform the same duties for the scientific community as a whole to be responsible, it also has some weaknesses. *Rescue #2* illustrates a case of "institutionalized" special responsibility. Lifeguards have acquired a special relation with people at risk of drowning. It is their role to save them and they have received special training to be able to perform it. It is in virtue of lifeguards' clear institutional role that the other people at the beach know that they ought not to try to rescue people in danger. In Rolin's account, however, the mechanism for the distribution of special roles across the members of a scientific community is not explained. Specifically, it is not clear whether the special relations that some individual scientists have somehow developed with other communities are sufficient to justify their having special roles.

One may imagine the case of a scientist who has developed a special relation with citizens but who does not necessarily do public engagement in a responsible way. Such a scientist may have gained popularity by using social networks and blogs, she has then increased her visibility by taking part in radio and TV shows, and so on. The fact that she has managed to develop a special relation with the public, however, does not necessarily make her a good fit to carry on the collective social responsibility of public engagement. This scientist may just present and explain data to the public in a biased way, pursue her own agenda, or make controversial claims that do not reflect the view of her scientific community. If this was indeed the case, would she be the only one to be held responsible for her wrongdoings, even though she was individually performing the collective duty of engaging with the public?

Rolin seems to put too much responsibility on some individual scientists (namely, those with special roles), while leaving the rest of the scientific community with too little responsibility. Despite her plea for considering the moral dimension of the scientific community, her view risks not accounting for the collective nature of the responsibility of the scientific community.

## 5. Scientific research groups as collectives

The existence of collective responsibilities does not imply that every individual member of a group is able or even required to carry them. Some collective duties ought to be carried by individuals who have received the “special role” of acting as a representative for their group. The risk with the distribution of special roles within a group is that collective responsibility could easily collapse onto individual responsibility, with the rest of the group being deresponsibilized.

As a solution to this order of problems, I propose to draw from the theories of philosophers such as French (1984), Rovane (1997), List and Pettit (2011), and Collins (2019). What all their theories have in the idea that collective responsibility can be attributed only to special kinds of organized groups.

As defined by Collins (2019), responsible collectives are “constituted by agents that are united under a rationally operated group-level decision-making procedure that has the potential to attend to moral considerations” (4). In her view, there are three necessary and jointly sufficient conditions for the establishment of a collective. First, each member must be committed (tacitly or explicitly) to act accordingly to the decisions of the collective. Second, the decision process of the collective considers the beliefs and preferences of its individual members, but it is “operationally distinct” from them (in the sense that it is not reducible to the aggregation of individual preferences). Third, the enactment of the collective decisions, made under a group-level procedure, requires individual members to perform duties and actions, where such individually carried duties and actions are attributable to the collective.

To clarify her view, Collins makes her own rescue example.

*Rescue #3* There are six strangers at the beach. One is drawing and the others are sunbathing . . . . The swimmer can be reached only with a motorboat. It will take two people to drag the boat to the water and hold it while a third strats it. The

boat will take off straight away, so the fourth and the fifth, who will pull the swimmer into the boat, must already be in the boat. Two are required to pull in the swimmer because none is strong enough to do so alone . . . . Each sunbather is wholly unknowledgeable about rescuing swimmers—except Laura . . . . Laura's instructions divide the necessary actions among the sunbathers. Each performs the action that Laura instructs him or her to perform, because Laura has instructed it. The swimmer ceases drowning. (Collins 2019, 108–9)

What differentiates Collins's and FS's example is the specification of how the random group *becomes* an organized mobilized group. Laura, who is the most knowledgeable person around, assigns different duties to different members of the group and checks that everyone is able to perform them. The other members commit to follow Laura's instructions and act on the belief that everyone will do the same.

The idea at the basis of *Rescue #3* is that because organized groups are more likely than random groups to produce morally desirable ends, random groups must first organize themselves in an optimal way. FS suggest that to prevent collective (epistemic) harms, individuals belonging to random collectives ought to “mobilize” the others. However, they do not discuss the problem of who or how many should join forces in forming a mobilized group, and then of how to organize the mobilized group once a sufficient number of people have joined forces. Following Collins, by contrast, collective responsibility cannot be ascribed to random collections of individuals if first they have not been “responsibilized” through the stipulation of an internal regulation for the distribution of tasks and duties.

Collins is aware that, *prima facie*, *Rescue #3* may look like an example of a “tyrannical” collective, in which one individual is in charge of the distribution of roles and duties, while all the others simply comply. To make her decisions, however, Laura receives the input of the other members of the group, who in turn may change the course of action in case Laura would show incompetence. For Collins, moreover, her theory of collectives can be extended to cover a wide range of groups of different sizes and levels of organization, from the basic level of groups of friends making collective decisions about dinner, to higher-level institutionalized groups, such as firms or corporations.

Finally, it is important to stress that, in Collins's framework, the assignment of special role-responsibilities to some individual members of the group is not the end of the story. The members of a collective who are not the carriers of special duties still have *membership duties*, that is, duties that “(i) are held by agents in virtue of the fact that they are members of a collective that has a duty and (ii) whose performance is a component of the collective's doing its duty” (Collins 2019, 181). Membership duties may be minimal. For example, they may amount to “first, checking to see whether the member that's supposed to be doing X is doing X, then, second, taking further action only if they're not” (*ibid.*, 197). Moreover, “the precise actions required will shift and change over time, becoming more or less demanding as the other members do (or don't do) their own membership duties. Also, one's duty will change as the content of one's role changes” (*ibid.*).

Although she develops a prescriptive moral framework for collectives “in general,” Collins suggests that her theory can work for specific groups, provided that one plugs in some relevant characteristics about them. My aim is to take a first step toward this

direction, by applying Collins's framework to a specific kind of group, namely scientific research teams.

Research teams have *both* epistemic and social responsibilities. These different responsibilities may be difficult to balance, as philosophers like Resnik and Elliott (2016) rightly point out. What could be problematic for an individual, however, may be harmonized through a collective approach.

Research teams already distribute different tasks to different individual members. Such a division of labor may also come with a distribution of epistemic responsibilities. Depending on who takes collective decisions within a research team and how some individuals (such as PIs and senior researchers) may carry more epistemic duties than others (such as junior researchers or lab technicians). A similar distribution of duties may work for social responsibilities too, which could be assigned to some individual members of the team, but not to everybody. In this way, some individuals will carry epistemic responsibilities, others will carry social responsibilities, or a mixture of epistemic and social responsibilities, while the research team as a collective will carry both epistemic and social responsibilities.

To work as a responsible collective, however, the members of a research team cannot simply load off social responsibility to one or few individuals. They have, in fact, membership duties, which require them to check whether the individuals with special roles are carrying their special responsibilities. Individuals with special roles, in turn, have to take into consideration the inputs of the other members of the team.

This is, I claim, how a responsible research group ought to be organized. This framework treats the responsibilities of science as collective responsibilities. In this way, it accounts for how a group of scientists may collectively carry different kinds of responsibilities without every single one of them having to carry them all. Furthermore, it avoids the problem of leaving individuals with special roles and special responsibilities unchecked or unsupervised.

## 6. Using the ideal, bringing it into the real

My normative theory of the collective responsibilities of science draws from some moral theories of collective responsibility and attempts to account for how epistemic and social responsibility may be harmonized at the group level. While this framework works "in theory," it still remains to see how to implement it in actual research groups. Of course, bridging the gap between a general normative theory and the actual scientific communities will require further work. Such a task will require, on the one hand, further theoretical articulation and, on the other hand, a finer understanding of how actual research teams reach collective decisions and distribute duties. Even at its present level of generality, however, my framework can be used as a "diagnostic tool" to bring to light and examine some issues in at least two different contexts.

The first is the context of the philosophical debate about values and responsibility in science. As already explained in section 3, such a debate is often conducted either from an individualistic perspective, which revolves around the responsible scientist, or from a more social perspective, which makes references to a rather generic scientific community but without tackling the issue of the individual variations within it. The approach of the philosophy of science to the discussion of the responsibilities of science needs to be reformed. Using the arguments and concepts

developed in moral philosophy to reframe the discourse about individual and collective responsibilities in science may be a first step. As shown in this article, in fact, moral philosophy provides some tools to distinguish different forms of responsibility (i.e., backward and forward looking, individual and collective) and to find a general definition of the social responsibility of science (which I have indeed defined as a form of collective remedial responsibility). Furthermore, moral philosophers have debated about the structure of responsible collectives; such debates may inform the discussion about the organization of well-designed responsible research groups.

My theory, moreover, can also serve to diagnose fallacies in how responsibilities are carried out by *actual* research groups. More and more often, scientists involved in projects funded by national and international funding agencies are demanded to comply with policies aimed at fostering social responsibility in research teams. Empirical evidence suggests that many researchers consider epistemic responsibility (i.e., the production of knowledge) as their primary if not only duty and regard the reflection on the societal impact of their work as “extraordinary”—that is, not common, not expected (Hennen et al., 2021). To comply with policies on socially responsible research, labs simply recruit ethicists or other scholars from the humanities and social sciences and leave them to “look after the ethics part” of specific projects. This easily leads to the delegation of social responsibility to few individuals (or even to one individual), rather than to its full integration within the whole research group (Delgado and Åm, 2018; Sigl et al., 2020). It has also been observed that research groups *already* implicitly assign some special duties—such as engaging with societal stakeholders and thinking about the potential implication of a project—to senior scientists and PIs, while junior scientists carry on the epistemic duty of producing knowledge through bench research (Politi 2024). As already discussed, delegating full responsibility to few individuals may deresponsibilize the rest of the group. Comparing the actual situation with an ideal model of responsible scientific collective may help questioning whether actual research teams display the *right* kind of distribution of duties, and beginning to ask what they should do in case they do not.

The problem then becomes: *How* can we make sure that actual research teams organize themselves in such a way that they display the right kind of distribution of duties? One concrete way to do so could be the establishment of research team policies. Elliott et al. (2017) argue that research teams and labs could create their own policies regarding coauthorship practices. Especially in large interdisciplinary collaborations, it may be difficult to pinpoint who or how many authors exactly contributed to a research article and an overinclusive approach to authorship may lead to questionable results. Without denying the importance of institutional rules, research teams could regulate themselves by establishing their own authorship policies. Similarly, labs and research teams could create their own policies incorporating plans for addressing social responsibilities. Such policies may be created at the beginning of a project and then revised as new team members join. They should also address the lack of power of junior researchers, who ought to carry out their membership duties even though they are not the carriers of some special duties. In short, lab policies would function like the *constitutions* defined by List and Pettit (2011). Constitutions play two functions: a *deterrent* function because they forbid the members of a collective from choosing some course of actions and a

*developmental* function because they regulate the performance of the members of the collective through time. The establishment of a policy/constitution would then be the first step toward the collective responsabilization of a team or lab.

Finally, when talking about the collective responsibilities of science, it is important to define the collective to which we are attributing responsibilities. Scientists form communities at different levels of generality: There is the community of all scientists, then there are “field” or “specialty-communities” (the biologists, the physicists, the sociologists, etc.) and “sub-specialty-communities” (the microbiologists, the quantum physicists, etc.). At the lowest level of generality there are research institutes, laboratories, and research teams. Rolin (2017) acknowledges the existence of scientific communities of different sizes and develops her framework for wide specialty-communities. Under this respect, Fleisher and Šešelja (2022) are less specific. However, the kind of harm they talk about may affect communities of specialists that are wider than small research groups; therefore, it is plausible to assume that their account also applies to specialty-communities.

By contrast, my account applies, first and foremost, to more “local” or “micro” communities, such as research teams. Science policies such as ELSI or RRI, which aim at implementing social responsibility in scientific research, are designed for such groups. It is their members, therefore, who face the challenge of harmonizing epistemic and social responsibilities. Ultimately, if “[the] paradigm group in science is a research team led by a Principle Investigator, working in the same laboratory on a well-defined project” (Bird 2022, 73), and if such teams are the basic unit of scientific knowledge production (Wray 2007), then they are also the agent with epistemic and social responsibilities. Without underestimating the need for institutional reforms at the higher level of specialty communities, my approach mainly contributes to solving the specific challenges posed by different kinds of responsibilities in research teams. This does not exclude that it could be scaled up to account for wider scientific communities.

Apart from research teams complying with science policies, it is commonsensical to expect that, for example, the communities of climate science, virology, or personalized medicine have epistemic as well as social responsibilities. The wider the scientific community, however, the more pressing and more difficult to solve the issue of collective organization. Specialty-communities have many members who, furthermore, work in geographically distant areas. Moreover, as discussed in the case of FS’s account, when a collective involves a very high number of individuals it will face the problem of insignificant hands.

Specialty-communities call for a reflection on the kind of group-level decision-making procedure they require to be held collectively responsible. As discussed in section 5, a collective requires its members to abide by the results of such a procedure, even if only implicitly. While it would be difficult to imagine a whole specialty-community organizing to choose directly its representatives from all the specialists working around the world, perhaps a first step toward the organization of such communities into responsible collectives could be analogous to that proposed for research teams and pass through the establishment of a common constitution. In the STS literature, remarks about the possibility, or even the necessity, of establishing a “constitution” for socially and politically responsible science, for example in the context of environmental research, have been made by Latour (1993, 2004). More recently, in philosophy, Mantzavinos (2021) has discussed the idea of “constitutional



arrangements introduced deliberately and by cultivating an accompanying culture of [social responsibility]" (ibid., 256). A constitution shared by a specialty-community may establish the "ethos" that its members must respect. It may establish, in other words, the membership duties of a specialty-community. If not in the assignment of special roles and duties, the constitution may be used to individuate those who do not respect the community's ethos, who do not carry their membership duties and who, therefore, ought to be considered as not belonging to, and not representing, the community.

## 7. Conclusions

Scientists have different kinds of responsibilities and carrying them all may be too demanding. Drawing from moral philosophy and, in particular, from theories of collective responsibility, in this article I developed a philosophical framework for an ideal responsible scientific community. Its aim is to show how dilemmas faced by individual scientists may be solved through a fair distribution of special roles within their research groups.

My proposal is only one of the possible ways for discussing the collective responsibilities of science. Developing a normative framework for an ideal responsible scientific community may be fruitful for the design of actual responsible scientific communities. Future debates will hopefully shed more light on the collective dimension of the responsibility of science.

**Acknowledgments.** I thank the anonymous reviewers as well as an associate editor of *Philosophy of Science* for their insightful comments. Funding for this research has been provided by the Beatriu de Pinós/ Marie Skłodowska-Curie Action Cofund fellowship (grant number: 2020-BP-00196)

## References

- Bird, Alexander. 2010. "Social Knowing: The Social Sense of 'Scientific Knowledge.'" *Philosophical Perspectives* 24 (1):23–56. <https://doi.org/10.1111/j.1520-8583.2010.00184.x>
- Bird, Alexander. 2019. "The Aim of Belief and the Aim of Science." *Theoria* 34 (2):171–93. <https://doi.org/10.1387/theoria.19351>
- Bird, Alexander. 2022. *Knowing Science*. Oxford: Oxford University Press.
- Bratman, Michael. 1992. "Shared Cooperative Activity." *The Philosophical Review* 101 (2):327–41. <https://doi.org/10.2307/2185537>
- Bratman, Michael. 1993. "Shared Intentions." *Ethics* 104 (1):97–113. <http://www.jstor.org/stable/2381695>.
- Bratman, Michael. 2013. *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Collins, Stephanie. 2019. *Group Duties: Their Existence and Implications for Individuals*. Oxford: Oxford University Press.
- Coop, David. 2007. The Collective Moral Autonomy Thesis. *Journal of Social Philosophy* 38 (3):369–388.
- Corlett, Angelo. 2001. "Collective Moral Responsibility." *Journal of Social Philosophy* 32 (4):573–84. <https://doi.org/10.1111/0047-2786.00115>
- de Melo-Martín, Immaculada, and Kristen Intemann. 2023. "Socially Responsible Science: Exploring the Complexities." *European Journal for Philosophy of Science* 13:33. <https://doi.org/10.1007/s13194-023-00537-6>
- Delgado, Ana, and Heidrum Åm. 2018. "Experiments in Interdisciplinarity: Responsible Research and Innovation and the Public Good." *PLoS Biology* 16 (3):e2003921. <https://doi.org/10.1371/journal.pbio.2003921>
- Douglas, Heather. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh: Pittsburgh University Press.
- Douglas, Heather. 2021. "Scientific Freedom and Social Responsibility." In *Science, Freedom, and Democracy*, edited by Péter Hartl and Adam T. Tuboly, 68–87. London: Routledge.

- Elliott, Kevin, Isis Settles, Georgina Montgomery, Sheila Brassel, Kendra Spence Cheruvilil, and Patricia Soranno. 2017. "Honorary Authorship Practises in Environmental Science Teams: Structural and Cultural Causes and Solutions." *Accountability in Research* 24 (2):80–98. <https://doi.org/10.1080/08989621.2016.1251320>
- Fleisher, Will, and Dunja Šešelja. 2022. "Responsibility for Collective Epistemic Harms." *Philosophy of Science* 90 (1):1–20. <https://doi.org/10.1017/psa.2022.9>
- French, Peter. 1984. *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Gilbert, Margaret. 1990. "Walking Together: A Paradigmatic Social Phenomenon." *Midwest Studies in Philosophy* 15 (1):1–14. <https://doi.org/10.1111/j.1475-4975.1990.tb00202.x>
- Gilbert, Margaret. 2000. *Sociality and Responsibility*. Lanham, MD: Rowman and Littlefield.
- Gilbert, Margaret. 2013. *Joint Commitment*. Oxford: Oxford University Press.
- Goldman, Alvin. 1970. *A Theory of Human Action*. Princeton: Princeton University Press.
- Hennen, Leonhard, Julia Hahn, and Miltos Ladikas. 2021. "Competing Modes of Responsibility in Research Organisations: Insights from an International Comparison." *Science and Public Policy* 48 (1):54–65. <https://doi.org/10.1093/scipol/scaa057>
- Hindriks, Frank. 2019. "The Duty to Join Forces: When Individuals Lack Control." *Monist* 102 (2):204–20. <https://doi.org/10.1093/monist/onz006>
- Hindriks, Frank. 2022. "The Problem of Insignificant Hands." *Philosophical Studies* 179:829–854. <https://doi.org/10.1007/s11098-021-01696-z>
- Kagan, Shelley. 2011. "Do I Make a Difference?" *Philosophy and Public Affairs* 39 (2):105–41. <https://doi.org/10.1111/j.1088-4963.2011.01203.x>
- Kitcher, Philip. 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- Kitcher, Philip. 2011. *Science in a Democratic Society*. New York: Prometheus Books.
- Kourany, Janet. 2021. "Bacon's Promise." In *Science, Freedom, and Democracy*, edited by Péter Hartl and Adam T. Tuboly, 88–109. London: Routledge.
- Latour, Bruno. 1993. *We Have Never Been Modern*. Cambridge, MA: Harvard University Press.
- Latour, Bruno. 2004. *Politics of Nature: How to Bring the Sciences into Democracy*. Cambridge, MA: Harvard University Press.
- Lefevere, Merel, and Erik Schliesser. 2014. "Private epistemic Virtue, Public Vices: Moral Responsibility in the Policy Sciences." In *Experts and Consensus in Social Science*, edited Carlo Martini and Marcel Boumans, 275–95. London: Springer. [https://doi.org/10.1007/978-3-319-08551-7\\_14](https://doi.org/10.1007/978-3-319-08551-7_14)
- Lewis, Hywel. 1948. "Collective Responsibility." *Philosophy* 23 (84):3–18. <https://www.jstor.org/stable/3747383>
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Longino, Helen. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.
- Longino, Helen. 2002. *The Fate of Knowledge*. Princeton: Princeton University Press.
- Mantzavinos, Chrysostomos. 2021. "Institutions and Scientific Progress." *Philosophy of the Social Sciences* 51 (3):243–65. <https://doi.org/10.1177/0048393120978453>
- Mele, Alfred, and Steven Sverdlík. 1996. "Intention, Intentional Action, and Moral Responsibility." *Philosophical Studies* 82 (3):265–87. <https://doi.org/10.1007/BF00355310>
- Miller, Boaz. 2015. "Why (Some) Knowledge Is the Property of a Community and Possibly None of Its Members." *The Philosophical Quarterly* 65 (260):417–41. <https://doi.org/10.1093/pq/pqv025>
- Narveson, Jan. 2002. "Collective Responsibility." *Journal of Ethics* 6 (2):179–98. <https://doi.org/10.1023/A:1015823716891>
- Pielke, Roger. 2007. *The Honest Broker: Making Sense of Science in Policy and Politics*. Cambridge: Cambridge University Press.
- Politi, Vincenzo. 2024. "Who Ought to Look towards the Horizon? A Qualitative Study on the Collective Social Responsibility of Scientific Research." *European Journal for Philosophy of Science* 14:19. <https://doi.org/10.1007/s13194-024-00580-x>
- Politi, Vincenzo, and Alexei Grinbaum. 2020. "The Distribution of Ethical Labour in the Scientific Community." *Journal for Responsible Innovation* 7 (3):263–79. <https://doi.org/10.1080/23299460.2020.1724357>

- Ratti, Emanuele, and Federica Russo. 2024. "Science and Values: A Two-Way Direction." *European Journal for Philosophy of Science* 14:6. <https://doi.org/10.1007/s13194-024-00567-8>
- Resnik, David, and Kevin Elliott. 2016. "The Ethical Challenges of Socially Responsible Science." *Accountability in Research* 23 (1):31–46. <https://doi.org/10.1080%2F08989621.2014.1002608>
- Rolin, Kristina. 2017. "Scientific Community: A Moral Dimension." *Social Epistemology* 31 (5):468–83. <https://doi.org/10.1080/02691728.2017.1346722>
- Rovane, Carol. 1997. *Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton: Princeton University Press.
- Rudner, Richard. 1953. "The Scientist *qua* Scientist Makes Value Judgments." *Philosophy of Science* 20 (1):1–6. <http://www.jstor.org/stable/185617>
- Rudy-Hiller, Fernando. 2022. "The Epistemic Condition for Moral Responsibility." In *The Stanford Encyclopedia of Philosophy* (winter 2022 ed.), edited by Edward N. Zalta and Uri Nodelman. <https://plato.stanford.edu/archives/win2022/entries/moral-responsibility-epistemic/>
- Schroeder, Andrew. 2022. "Thinking about Values in Science: Ethical Versus Political Approaches." *Canadian Journal of Philosophy* 52 (3):246–55. <https://doi.org/10.1017/can.2020.41>
- Šešelja, Dunja, and Christian Straßer. 2014. "Heuristic Reevaluation of the Bacterial Hypothesis of Peptic Ulcer Disease in the 1950s." *Acta Biotheoretica* 62 (4):429–54. <https://doi.org/10.1007/s10441-014-9219-7>
- Sigl, Lisa, Ulrike Felt, and Fochler Maximilian. 2020. "I Am Primarily Paid for Publishing . . .": The Narrative Framing of Societal Responsibilities in Academic Life Science Research." *Science and Engineering Ethics* 26:1569–93. <https://doi.org/10.1007/s11948-020-00191-8>
- Sinnott-Armstrong, Walter. 2005. "It's Not My Fault: Global Warming and Individual Moral Obligations." In *Perspectives on Climate Change: Sciences, Economics, Politics, Ethics*, edited by Walter Sinnott-Armstrong and Richard B. Howarth, 293–315. London: Elsevier. <https://doi.org/10.1093/oso/9780195399622.003.0029>
- Smiley, Marion. 2022. "Collective Responsibility." In *The Stanford Encyclopedia of Philosophy* (fall 2023 ed.), edited by Edward N. Zalta and Uri Nodelman. <https://plato.stanford.edu/archives/fall2023/entries/collective-responsibility/>
- Sverdlik, Steven. 1987. "Collective Responsibility." *Philosophical Studies* 51 (1):61–76. <https://doi.org/10.1007/BF00353963>
- Talbert, Matthew. 2023. "Moral Responsibility." In *The Stanford Encyclopedia of Philosophy* (fall 2023 ed.), edited by Edward N. Zalta and Uri Nodelman. <https://plato.stanford.edu/archives/fall2023/entries/moral-responsibility/>
- Timmermans, Job, and Vincent Blok. 2021. "A Critical Hermeneutic Reflection on the Paradigm-Level Assumptions Underlying Responsible Innovation." *Synthese* 198 (suppl 19):4635–66. <https://doi.org/10.1007/s11229-018-1839-z>
- van der Burg, Simone. 2009. "Taking the 'Soft Impacts' of Technology into Account: Broadening the Discourse in Research Practice." *Social Epistemology* 23 (3–4):301–16. <https://doi.org/10.1080/02691720903364191>
- Watkins, John. 1957. Historical Explanations in the Social Sciences. *British Journal for the Philosophy of Science* 8 (30):104–117.
- Wilholt, Torsten. 2016. "Collaborative Research, Scientific Communities, and the Social Diffusion of Trustworthiness." In *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*, edited by Michael Brady and Miranda Fricker, 218–34. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198759645.003.0012>
- Wray, Brad K. 2007. "Who Has Scientific Knowledge?" *Social Epistemology* 21 (3):337–47. <https://doi.org/10.1080/02691720701674288>