

Who ought to look towards the horizon? A qualitative study on the collective social responsibility of scientific research

Pre-print version of the article published in *European Journal for Philosophy of Science*, 2024, DOI: <https://doi.org/10.1007/s13194-024-00580-x>

Vincenzo Politi

vin.politi@gmail.com

ORCID: <https://orcid.org/0000-0002-5753-3196>

Abstract. There is a growing concern for the proper role of science within democratic societies, which has led to the development of new science policies for the implementation of social responsibility in research. Although the very expression ‘social responsibility of science’ may be interpreted in different ways, many of these emerging policy frameworks define it, at least in part, as a form of anticipative reflection about the potential impacts of research in society. What remains a rather under-discussed issue is the definition of the bearer of the social responsibility of science. In other words, it is not clear who is supposed to engage in such an anticipative reflection, whether individual researchers or research groups. In the past few years, philosophers of science have begun to use qualitative research methods to fill the gaps between normative models of the organisation of ideal scientific communities and the reality of actual scientific practices. In this article, I follow this approach to discuss the issue of the collective dimension of the social responsibility of science. I rely on a qualitative study conducted on an interdisciplinary research group and I describe how group dynamics position individuals and distribute duties and roles, including social responsibility. Qualitative descriptions of the distribution of duties within actual research groups should inform the formulation of general prescriptive theories on the collective responsibility of science.

Keywords: Social responsibility of science; Collective responsibility; Qualitative methods in philosophy of science; Scientific practice; Social Epistemology of Responsible Science

1. Introduction

Scientists must produce reliable knowledge and avoid the spread of false information: they have, in other words, epistemic responsibilities (Kornblith 1983). They also have other responsibilities. For instance, they must not harm society. For this reason, committees and professional integrity codes, both at the national and at the international level, regulate and put legal and ethical constraints to research.

However, scientists also have a kind of responsibility that is different from the production of reliable knowledge and harm prevention, and which goes under the name of social responsibility. Not only scientists are expected to ‘do the right things’, by not causing harm and by respecting integrity codes, they also have to align their research to the needs of society and to democratically held values.

In the past few decades, debates about the social responsibility of science have become prominent among philosophers and STS scholars (Longino 1990, 2002; Funtowicz and Ravetz 1993; Gibbons et al. 1994; Intemann and de Melo-Martín 2010; Intemann 2015; Kitcher 2001, 2011; Douglas 2009; Kourany 2010). They are also central in science policy. For example, one of the work programs of the EU 7-year funding scheme Horizon2020 was ‘Science with and for Society’ (*SwafS*) (European Commission 2014; Figueiredo Nascimento et al. 2016). Among other things, this work program supported projects devoted to the development and institutionalisation of a new science policy, called ‘Responsible Research and Innovation’ (RRI) (Von Schomberg 2013; Owen, Stilgoe, Macnaghten 2013). RRI and similar policies are not limited to the introduction of new top-down structural and institutional rules, such as the requirement of publishing Open Access, of engaging with societal stakeholders, of finding new avenues for dissemination of results, and so on. They also aim at bringing forth a shift in research culture through a series of interventions targeting the so-called ‘mid-stream’ stage of science and involving researchers themselves (Fisher, Mahajan, & Mitcham 2006). This is motivated by the idea that, for science to be socially responsible, scientists ought not be passive recipients of other societal actors’ needs and perspectives. Rather, they ought to learn how to discuss and assess social concerns explicitly, and to reflect upon their own moral and social stances. To say it like feminist philosopher Donna Haraway, responsibility requires the ability to respond, or ‘response-ability’ (Haraway 2016:130).

In this article, I address a rather under-discussed philosophical problem, namely the definition of the bearer of social responsibility in science. Scientific research is carried out by different individuals, with different skills, attitudes, and perspectives. Perhaps it would be unrealistic to expect from every single one of them the same level of social responsibility. One may even wonder whether, for a research group as a whole to be socially responsible, it is even necessary that every single one of its members carries the duty of being socially responsible in the same way, at the same time, and with the same degree. A normative theory of the social responsibility of science should specify whether such a responsibility is individual or collective; and, in the latter case, it should specify which kind of group structure ensures that science can be collectively responsible.

Instead of proceeding in an armchair fashion, I propose that such a normative theory should be grounded on, and developed from, the description and analysis of actual research groups. This is why, as a first step towards the development of a theory of the social responsibility of science, in this article I present the results of a qualitative study on an interdisciplinary research group complying with RRI. The research questions driving the study were: Who are the ‘socially responsible’ members of the research group? Does their engagement in social responsibility make the whole group socially responsible? Is the research group a collective agent?

After defining the social responsibility of science as a form of *collective remedial responsibility* (sections 2-3), I discuss the results of the case study (sections 4-6), from which I will draw some philosophical lessons concerning the distribution of duties within research groups (section 7). Finally, I will return to some methodological issues to explain how my specific case study works as a first step towards a more general *social epistemology of responsible science* (section 8).

2. Defining the social responsibility of science

‘Social responsibility of science’ is a rather polysemous expression (Timmermans and Block 2021; de Melo-Martín and Intemann 2023). Nevertheless, it is possible that the different definitions of the social responsibility of science share some commonalities. For instance, the need to anticipate the potential implications of research at the social and cultural level is central not only to RRI (Fisher and Rip 2013; Owen et al 2013; Von Schomberg 2013), but also to frameworks such as Technology Assessment (TA) or the Ethic, Legal, and Social Implication research in innovative sciences (ELSI).

In short, many science policies that seek the implementation of social responsibility in scientific research conceptualise it as including a form of anticipative reflection as one of its central components.

Moral philosophers distinguish between *backward-looking* and *forward-looking responsibility*. So-called 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, instead, is the moral obligation to begin a course of action to improve the current situation. It can be either *preventionist* (responsibility as avoiding future harms) or *remedial* (responsibility as the commitment to bring forth positive changes). Following a famous example made by Parfit (1984), an agent who is responsible in a preventionist sense avoids leaving a piece of broken glass in the park in order to avoid harming an unfortunate child in the future. ‘Harm’, in a sense, is time indifferent: what counts as harmful today would count as harmful anytime. This is why, qualitatively speaking, the future of preventionist responsibility is an extension of the present. What Parfit does not consider is remedial responsibility that amounts to reflecting upon values (to use his own example, reflecting upon child safety and environmental protection), considering necessary transformations (like putting baskets for separate waste collection in the park), and acting accordingly (by starting a petition to ask the municipality to install such baskets in the park). The outcome of the responsible action may even alter the standards for assessing what is valuable. Because of the responsible action of our example, the landscape of the park will be modified, or the closing hour will be anticipated to allow for the collection and segregation of waste. In the future, however, the beauty of the park and its accessibility could be considered less important than they are today.

General accounts of responsibility typically involve both backward- and forward-looking aspects. The emerging science policies mentioned in this article are not concerned with the mechanisms for blaming or punishing scientists for misconducts or wrongdoings in an after-the-fact fashion. Rather, they aim at promoting the cultivation of anticipative reflection among researchers, with the idea that such a form of reflection is necessary to make science socially responsible. This is not to say, of course, that RRI and similar policies are ‘dismissing’ the crucial importance of backward-looking responsibility, but only that they aim at filling a gap in the governance of science by attempting to institutionalise forward-looking responsibility.

That science ought not to be harmful, moreover, is commonsensical. Besides, committees and professional integrity codes are already in place to prevent harm, by regulating and putting legal and ethical constraints to research. The potential implications of scientific and technological research, however, are not always clear-cut and somehow quantifiable harms (as in the case, for example, of technological devices that may clearly compromise the health of end-users). Science and technology, in fact, may also involve so-called ‘soft impacts’ (as in the case of technologies transforming how people interact, socialise, and perceive themselves; see van der Burg 2009). Not only may they have unintended practical consequences, but science and technology may also transform a currently held value system in unintended and unexpected ways (Ratti and Russo, 2024). Although necessary, quantitative risk/benefit analyses for preventing harms may not be sufficient for socially responsible research, which also requires a critical reflection on the transformative and disruptive power of innovation (Jasanoff 2016).

From these considerations, it follows that the social responsibility of science that the newly emerging policies seek to implement within research groups is a form of *forward-looking remedial responsibility*. As such, it requires the ability of imagining the transformative future impacts of science and technology.

Finally, it is important to mention that not everybody agrees with the idea that social responsibility involves anticipative reflection. Both Nordman (2014) and Carrier (2021) discuss the impossibility of knowing the future pathways of scientific and technological developments. Even though, to some extent, it is possible to anticipate what could happen in the world as we know it, science and technology may transform and reshape reality in unexpected ways. Carrier and Nordman point out that while the unpredictability of innovative science and technology is precisely the reason that has led to the need of formulating new policies, such policies end up requiring the impossible, that is predicting the future impacts of innovative research. As several commentators argue, however, the issue here is what one means by ‘anticipating the future’ and what one expects from such an activity (Selin 2014; van der Burg 2014; Wilsdon 2014; Urueña 2021). Obviously, anticipative reflection is not an exercise in ‘fortune telling’. Cultivating imagination about desirable future scenarios and how to bring them about, however, may drive research towards specific directions, while also taking into consideration social values. That the future is unknowable as a matter of fact does not

preclude scientists to be ‘oriented’ towards some possible future scenarios rather than others.

3. Collective responsibility

When we talk about the social responsibility of science, we are not talking about the responsibility of any one scientist in particular. This is so because scientific research is not (and cannot be) carried on by lone individuals, but by groups of collaborating researchers. Moreover, as argued by Wilholt (2016), the public trust on science is not directed towards individual researchers, but rather towards collective bodies, which are therefore to be considered as the responsible agents of science. This means that the social responsibility of science, which in the previous section has been defined as a kind of remedial forward-looking responsibility, should also be defined as a kind of *collective responsibility*.

The concept of collective responsibility is highly debated in moral and political philosophy (Smiley 2022). To be responsible, a moral agent ought to meet (at least) three conditions: the *epistemic condition* (possessing the right kind of knowledge and awareness), the *freedom condition* (the ability to act free from external constraints), and the *intentionality condition* (responsible agents must intend to do the right thing, meaning that responsibility is not attributed to someone who ends up doing the right thing unintentionally, let alone out of sheer luck). Supporters of the concept of collective responsibility argue that, in some cases, responsibility can be attributed not only to individual people, but also to groups of people. The idea that a group may know more than the sum of the pieces of knowledge possessed by its individual members, and that a group may do and accomplish things that individuals alone would not be capable of, lay the ground for the viability of the concept of collective responsibility.

One of the main issues is whether groups also possess agency. Traditional arguments against collective responsibility reject the existence of collective intentions and affections, necessary for the ascription of moral agency to groups (Lewis 1948; Watkins 1957; Goldman 1970; Sverdlik 1987; Corlett 2001; Naverson 2002). Contemporary defenders of collective responsibility, however, do not make controversial metaphysical claims about the existence of a ‘collective mind’. In their views, collective agency does not require the existence of such a mysterious entity, but only of specific group structures that facilitate the so-called ‘we-reasoning’, that is distinct from the aggregation of many ‘I-reasonings’ (Schwenkenbecher 2019).

For example, Gilbert (1987, 2000, 2013) argues that a *plural subject* is established when different individuals are unified by a ‘joint commitment’, that is when they all have the same objective and act accordingly. For Bratman (1992, 1993, 2013), collective intentionality requires something more than joint commitment. Individuals acting in the same way towards the same goal may lack mutual communication and cooperation: in such a case, they would be united by a joint commitment but disjointed under many other respects. The group they form, therefore, could hardly be considered as a collective agent. This is why, for Bratman’s, collective intentionality can be attributed only to *collective subjects*, whose members are engaged in a ‘shared cooperative activities’.

Other philosophers, like French (1984), Rovane (1988), List and Pettit (2011), and Collins (2019) argue not any group displaying shared intentionality and coordination can be collectively responsible. For example, Collins (2019) develops a ‘tripartite model’ of groups, that consists of: *collectives*, “constituted by agents that are united under a rationally operated group-level decision-making procedure that has the potential to attend to moral considerations”; *coalitions*, formed by “agents who each hold a particular goal and are disposed to work with the others to realise the goal, while lacking a group-level decision-making procedure that has the potential to attend to moral considerations”; and *combinations*, formed by “any collection of agents that do not together constitute either a collective or a coalition” (Collins 2019:4). In her view, collective responsibility can be attributed only to collectives. Collins expands on the literature on so-called *corporate responsibility* (French 1984, Rovane 1988, List and Pettit 2011).

Without entering in too many details, Collins and other authors focus on the level of internal organisation a group ought to display in order to be attributed collective responsibility. For example, collectives may distribute different duties to different individuals. Another characteristic of collectives is not only their ability to reach a goal, but also to set new goals and, in case, even to change them.

On the one hand, the idea that scientific communities can be regarded as collective agents is becoming increasingly popular in philosophy of science. Social epistemologists have even developed formal models of the social organisation of science in order to solve the “mismatch between the demands of individual rationality and those of collective (or community) rationality” (Kicher 1990:6). The problem is that the debate in social epistemology seems to be limited to the consideration of the

scientific community as a collective *epistemic* agent. For instance, little or no attention is paid to the potential mismatch between the demands of individual responsibility and those of collective responsibility.

On the other hand, the philosophical debate about the responsibility of science often revolves around the analysis of the roles of non-epistemic values in research, but with no further specification of *who*, within a research group, holds such values. Saying that ‘the scientist *qua* scientist makes value judgements’ (Rudner 1953) may amount to a discussion of ‘the responsible scientist’, a rather idealistic individual that represents no actual scientist. Philosophers who focus on the communitarian dimension of the social responsibility of science (i.e., Longino 1990, 2002; Kitcher 2001, 2011), by contrast, often end up talking about ‘the responsible scientific community’, which is no less idealised than its individualistic counterpart and which does not help understanding how different individuals, with different attitudes and skills, may contribute to group-level decisions.

In short, while social epistemology misses a ‘moral dimension’, the philosophical debate about values and responsibilities in science appears to miss a ‘social dimension’. This gap could be filled by a *social epistemology of responsible science*. In this paper, I take one of the first steps towards such a project by reporting the results of a qualitative study on an actual research group and by then analysing whether such a group does possess collective moral agency.

4. The case study

I will present and discuss some of the results of a qualitative case study to analyse: whether and how the individuals of a research group engage with anticipative reflection, or fail to do so; what kind of potential future impacts they are able to anticipate and reflect upon; and how the duties of thinking about the future impacts of research are (implicitly) distributed within the group.

The study has been conducted on PINpOINT (PP), a 3-year project in precision medicine at the University of Oslo (UiO, Norway). The aim of PP is to identify the most effective individual treatments for some types of blood cancers (specifically, multiple myeloma and some forms of leukaemia) and to develop an AI-based diagnostic tool capable of forecasting how specific treatments would work in different individuals.

The PP team includes clinicians from the Department of Haematology at the Oslo University Hospital (OUH), molecular biologists and bioinformaticians from the Institute for Cancer Research (ICR), and biostatisticians from the Oslo Centre for Biostatistics and Epidemiology (OCBE). The PP working pipeline unfolds as follows:

- The clinicians from OUH extract genetic material (cancer cells and bone marrow) from consenting patients, and send it to ICR.
- At ICR, molecular biologists specialised in multiple myeloma and in various forms of leukaemia run *ex-vivo* experiments, consisting in administering different drug combinations in varying doses to the genetic material. Considering the number of possible drug combinations and dose variations, the amount of data generated after each experiment is huge. The data are first visualised and ‘cleaned’ by bioinformaticians at ICR, and then sent to the biostatisticians.
- At OCBE, one group of biostatisticians works on *statistical models* to estimate the so-called ‘synergistic effect’ (Berenbaum 1989; Greco, Bravo, and Parson 1995) of the drug combinations from the experiments. Another group develops an AI-based *mechanistic model* (Barbolosi et al 2016) to forecast the evolution of specific drug therapies on individual patients.

The pipeline is in its experimental phase, meaning that the experiments and data analyses are not concluded, and that the mechanistic model is not fully developed yet. Moreover, for the results of PP to be fully validated, after the experimental phase it will be necessary to make *in vivo* experiments on animal models and, eventually, to run the appropriate trials on human subjects.

PP is supported by the Centre for Digital Life Norway (DLN), a national centre which facilitates innovative projects combining biotechnologies with digital technologies in fields such as health, aquaculture, and agriculture. DLN is in turn funded by the Research Council of Norway (RCN), which requires each grant proposal to include an account of how the research is conducted according to RRI. As conceptualised by DLN, RRI aims at fostering a scientific culture where “anticipating and reflecting upon the known, as well as unknown, desirable as well as undesirable

impacts of said research is the norm rather than the exception” (<https://www.digitallifenorway.org/competence-areas/rri/>). In complying with this requirement, PP includes an RRI work-package (which the present study is a part of).

5. Methodology

From September 2020 to March 2022, I conducted a qualitative study on the members of the PP research group to uncover their level of anticipative reflection for the potential impacts of their project. The sources of data have been: field observations and guided visits at the laboratories; participation to the monthly PP meetings; one-on-one meetings (either in person or on Zoom) with some of the PP members, who explained in more details the research conducted by the various sub-teams of the pipeline; document analysis (articles written by PP members, the PP project proposal, the work packages plan, and so on); field notes and a personal journal; the organisation and participation of a public engagement event; and qualitative semi-structured interviews on some members of PP.

Although in the present article I mainly discuss the results of the interviews, it is important to stress that the other sources of data were necessary for a number of reasons. As a philosopher entering an interdisciplinary research group working on precision oncology, it was crucial for me to socialise with the other members of PP. The socialisation process allowed me, on the one hand, to gain some level of trust with the scientists and, on the other, to acquire at least a very basic understanding of the methods, aims, and theoretical vocabulary of PP. Socialisation through participant observation, therefore, was necessary to formulate the right kind of questionnaire and to understand the answers of the interviewees. Moreover, field observations, document analysis, and other activities were not conducted only before the development of the questionnaire and did not end with the interviews. These activities were conducted in parallel with the qualitative interviews and they were used to triangulate my findings and help refine my analysis.

The sample for the interviews consisted of 13 PP members. Of these, 8 work at ICR (4 of them belong to the sub-group studying multiple myeloma, 4 to the sub-group studying leukaemia) and 5 work at OCBE (2 of them belong to the sub-group developing statistical models of drug synergy effect, 3 to the sub-group developing the AI-based mechanistic predictive model). As for the level of seniority of the participants, 3 of them are principal investigators, 6 are permanent researchers at UiO, 2 are

postdoctoral researchers, 1 is a PhD student and 1 is an MA student who has recently obtained her degree and who is going to start a PhD within PP. The international research culture at the University of Oslo is also reflected in the sample: 5 of the interviewees are from Norway, while 8 are from abroad. The male/female ratio is 60-40%. Participants were suggested by the project PIs, who picked up volunteering members of their subteams. Interviews were carried in English, either face-to-face or on Zoom; in the first case, at the interviewees' workplace (i.e., office or laboratory), in the second case from their homes. Apart from me and the participant, no one else was present at the time of the interview.

The interviews were recorded, anonymised, and then thematically analysed with the *N-Vivo* software and by following the method of *Grounded Theory*, which involves the application of inductive reasoning to derive themes from data interpretation (Chamaz 2014). This means that the data collected through the interviews were not used to test a hypothesis but, rather, to develop new ones. 'Inward' and 'outward impacts', the way they relate to epistemic or social responsibility, and all the other concepts discussed in the next section were not formulated before this study, but *through* this study. Whenever this article directly cites some of the interviewees, pseudonyms have been used. I presented and discussed the results of this study in some of the PP monthly meetings, to make sure that all the participants felt comfortable with how the study was presented and to further consolidate my interpretation of the data.

Even though this study was not carried on with the specific aim of stimulating a *change* in researchers' anticipative reflection, so-called 'reactivity' cannot be excluded, especially considering that, during the interview, participants were asked to focus on, and start thinking about, the societal impacts of their own work. Clear cases of reactivity were indeed observed. For example, 'Charlotte' (senior researcher) claimed that, as she was answering the questions, she was beginning to think about modifying her protocol to take into consideration new parameters which may facilitate the future application of her results. 'Emily' (senior researcher) discussed with me how she began thinking about how to contribute to the implementation of the results of parts of PP. (Since both Charlotte and Emily spoke about technical aspects of their ongoing research, what they said will not be reported here.) These are only the most striking observed examples of shifts in the interviewees' reflection about their own work. Other 'micro-shifts' were also observed. In fact, one of the main themes that emerged in the data analysis, the '*Good Question*-theme', is about all the instances in which researchers

answered by exclaiming “*This is a good question!*”, or with similar expressions. Such exclamations were followed by moments in which researchers appeared to reflect about particular questions or issues.

In general, while it must be acknowledged, reactivity in qualitative studies does not necessarily affect the quality of data (Zahle 2023). In this research in particular, that some answers or behaviour were interpreted as being caused by an *individual* shift in anticipative reflection is not particularly problematic. The aim of this study, it is worth reminding, is to study how anticipative reflection is distributed across the research *group*. At least in the short term, individual increasing in anticipative reflection does not alter its distribution at the group level. In any case, unfortunately repeat interviews could not be carried on, due to the end of my appointment at UiO.

6. Results

6.1 Intentions

Seven PP members are driven by a preoccupation with contributing to society and they appear to have chosen a career in scientific research because they are driven by the desire of bringing positive social changes. They regard science as embedded into a much larger social context and they see the attempt of making a valuable contribution as one of their chief duties. This does not mean, of course, that one can infer that the other six interviewees do *not* see their research as aiming to benefit society. It may be the case that they are also driven by similar motivations, but they have not expressed them explicitly. The *Socially-Oriented*-theme, in fact, emerged inductively from the answers to general questions about their career path or their role within PP.

Julie (mid-level researcher): “I like the translational character of this field. It is more connected to patients, to really delivering something that society or the public can [use]... I would not say it’s ‘more useful’, I think it is just more translational”.

Mark (PI): “I have been interested in the ethical aspects [of science] all my life [...]. I have always been interested in understanding the long term effects of science, and what to do”.

Angel (postdoc): “I wanted to do something that had a more direct impact. I wanted to work on research that has a more immediate impact”.

While one may expect an interest for translationality from biomedical researchers working to fight cancer, biostatisticians, who perhaps could be more easily regarded as interested in analysing numerical data rather than in ‘real world’ problems, also show a strong desire to contribute to society. One biostatistician recalled how scientific research and political activism have been strongly intertwined in his life since his years as an undergraduate studying pure mathematics. Another one explained that, like many graduates in statistics, he could have opted for a more lucrative career in fields such as finance. Instead, he deliberately chose a different path because he values the potential of contributing to the well-being of people more than personal financial gains.

As mentioned in section 3, one of the conditions for the attribution of responsibility is intention: an agent is responsible for something if she has (or has had), among other things, the intention to do that. However, it is worth stressing that intention is only *one* of the necessary conditions for responsibility. It is therefore crucial to understand whether PP researchers are also able to translate their positive intentions, motivations, and social orientation into a socially responsible way of approaching their scientific activity.¹

6.2 Impacts and futures

To understand whether they engage with anticipative reflection, and of which kind, PP researchers were asked a number of questions about the potential impacts of the

¹ An anonymous referee, whom I thank, notices a tension between having a developed theoretical framework for analysing responsibility and then pursuing a grounded theory approach. For example, the referee argues, the relevance of intentions is given by the theoretical framework and then it is used to analyse the data. Two things must be stressed. First, the *Socially Oriented* theme emerged inductively from the answers given to general questions about the researchers’ scientific trajectory or about their role at PP, at what could be defined as the ‘first level’ of analysis. None of the questions mentioned terms such as ‘intention’, ‘responsibility’, ‘ethics’, ‘social orientation’, and the likes. It is only at a ‘second level’ of analysis that the *Socially Oriented* theme is interpreted as akin to one of the necessary conditions of responsibility. In other words, intention *became* relevant not because of a theoretical framework about responsibility, but because of the data. Second, even the second level of data analysis does *not* involve the interpretation of data through a developed theoretical framework. By definition, a responsible agent must meet some conditions, one of which is intentionality. Such a basic definition does not correspond to a fully developed theory of responsibility.

results of their research. The answers suggest that PP researchers reflect upon two different kinds of impact, that I associate with two different kinds of responsibility.

Eight researchers anticipated what I define as *inward impact*, that is the impact that scientific results may have on scientific research itself. Inward impact is talked about in a matter-of-fact fashion: researchers assume that they will obtain some results that the scientific community will eventually find useful and applicable. They also know which kind of results they need to produce, and how, in order to make a relevant inward impact. Even though they cannot know exactly when such results will be reached, researchers expect to make an inward impact sometimes in the medium-term future, that is, by the end of PP or of a follow-up project. The medium-term future researchers anticipate is not qualitatively different from the present: the kind of knowledge produced today will still be regarded as a valuable piece of knowledge in the future.

Inward impact is the result of present efforts, it amounts to the production of shareable knowledge the value of which, in the future, will be measured by current standards. For this reason, I associate inward impact with *epistemic responsibility*. Internal impact is made only if researchers act in an epistemically responsible way towards their colleagues and peers: that is, only if they produce and share robust, scientifically sound, and non-trivial results.

Outward impact, by contrast, is the impact a research may have in the world outside the laboratories and the specialists community. It belongs, in short, to the kind of societal implications that emerging science policies such as RRI require researchers to anticipate. A total of eight researchers talked about outward impact. However, references to outward impact are not only less frequent, but also characterised by a sort of ‘wishful thinking’. Researchers ‘hope’ that their work will make an outward impact (i.e., that it will be useful, that it will improve the quality of life of cancer patients, and so on), but they are not able to anticipate whether, when, how, and by whom such an impact will be made.

Alice (master student): “[PP] will have an impact *of course* on other researchers who are doing the same work like us, because we provide new knowledge. And *hopefully* it can also impact on how we, in the future, decide how to treat the patients: if we can show that our approach can somehow predict the response in the patients, then in the longer term maybe we can implement it in the clinic and can have an impact for the

treatment of the patients. But then it's... But then this will take several years and we have to show it in clinical trial studies that it is actually working, that it can predict the treatment response, and also we need to communicate with the clinicians, so that they have trust on what we are doing so they want to use it also”.

Anthony (mid-level researcher): “[PP can affect] the [research] community, us as a project first. It has an impact on the community at large, like the international community [of researchers]. The question we are trying to answer could be useful to others to build on. And of course we are aiming to make an impact on patients, right? So, we want to improve the way patients are treated at the moment. That's the ultimate goal, I would say, for this kind of project, and whether you will reach it, with this project or with some others later on, it's difficult to say. It's very difficult to say yet. But we have *hopes*, of course”.

I: *Who do you think this research project will affect?*

Sten (PhD student): “[*Sighs*] Who is going to benefit from it? I mean, the *hope* is always that we are able to deliver better treatment, right?”

I: *Who do you think this project is going to have an impact on?*

Emily (senior researcher): “I *hope* cancer patients will get benefits from it, *at some point*”.

Outward impact, if any, is achieved in a vague long-term future. The more distant and vague it becomes, the less researchers see themselves as agents capable of contributing to it. This suggests that many researchers do not see themselves as playing a prominent role in making an outward impact. Some of them explicitly claim, for example, that the successful implementation of the results of PP depends on the clinicians and on their willingness to adopt an AI-based tool they may not fully understand or even trust. Yet, they do not always see themselves as being responsible to contribute to clinicians’ awareness, appreciation, and understanding of such an innovative healthcare tool. Others consider that the implementation of innovative

healthcare tools may face financial barriers, which need to be tackled by politicians, by the National Health System, and maybe by patient associations as well. In short, it appears that PP researchers tend to delegate the responsibility for making an outward impact to other societal actors, despite the fact that, as seen in 6.1, many of them claim that the possibility of contributing to society is one of their main motivations for pursuing scientific research.

It must be said that several interviewees do *not* exclusively think about only one kind of impact. When asked to consider ‘consequences’, ‘impacts’, ‘implications’, ‘potential barriers to implementation’ and ‘who will be affected’, the answers of four researchers clearly referred to a mix of inward and outward impact. What is important to stress is how differently these two impacts are anticipated and assessed. They are in fact imagined as happening in two different kinds of future: a well-defined medium-term future for inward impact, in which researchers see themselves as playing an active role, and a rather vague and undefined long-term future for outward impact, for which researchers feel less responsible. (Table 1 summarises the relations of the concepts emerged from the data).

6.3 Deviations

Four individuals stood out from the rest of the PP research group.

One of them, Thomas (postdoc), talked about *personal impact*, that is, the impact research may have on one’s own future as a researcher.

Thomas (post-doc): “Of course I’m worried about [things going wrong with the project], because my academic livelihood depends on it. [...] My success in being able to finish my PhD depended on whether I could finish that (project), and I was worried about it every day. And now that I have a PhD, whenever I get something which I do not expect from the simulation I still panic, and I reach out to check whether I am correct or not”.

Thomas considers the short-term future of the project, because he is worried about getting wrong results, as well as his own future in research, because he is worried about

the effects of errors or dead-ends for his academic career. While Thomas *also* discusses the internal and the external impact of PP, he is the only one among the interviewees to talk openly about personal impact. However, it cannot be excluded that the others also consider personal impact simply because they did not talk about it explicitly.

The other three individuals are PIs. ‘Mark’ is a *questioning researcher*. Like the well-intentioned researchers discussed in 6.1, he is driven and motivated by the desire of bringing a positive change into society, and has chosen a career in academic research accordingly. What distinguishes him from the others is that his firm commitment to social and political values leads him to be deeply self-reflective and to question his own activity. The questioning researcher does not have answers, because he is aware of the complexity of the ethical and social issues arising in innovative biomedical research. Such a critical (and self-critical) reflection, however, does not appear to translate into a change on how research is conducted. When talking about the role of scientists in society, and despite his commitment to social and political values, the questioning researcher actually seems to support a rather traditional conception of science. In his view, scientists have to provide data to policy-makers, who will then have the responsibility of choosing what to do.

I: *Apart from giving results and data, do you also give any suggestions? Any policy suggestions? Why not?*

Mark (PI): “No, because this is another type of... So, we describe the situation now and we make a prediction of how it will continue in [the near future], if nothing is changing... But to ask the question ‘What would happen if I [did this]?’ is something that I cannot do in my models, because I don’t have that macroscopic... It’s a little bit like, ‘What would happen in my model if I change the gene?’, I cannot answer this question, because I do have that gene in my model”.

Claims like this show how complex forward-looking remedial responsibility is. The questioning researcher is indeed driven by moral and social values, but he also believes that the best way for serving social aims is the production of rather ‘value-free’ models, without engaging in the kind of future-oriented counterfactual reasoning necessary to the kind of social responsibility promoted by some of the emerging science policy

frameworks. Taking a step back from the abstract ideals of science, it is possible to see a more complicated reality in which different kinds of scientists occupy different places in the varied spectrum between value-freeness and value-ladenness.

As said above, most PP-researchers think only about the potential inward impacts of the research outcome of their own sub-team, but do not seem to have a bird-eye view of PP as a whole. Such a bird-eye view can be however found among senior researchers, especially PIs. This is a rather expected result: since they designed it to begin with, it does not come as a surprise that PIs are also able to look at the project pipeline in its wholeness. ‘Matthew’ is defined here as a *forward-looking researcher* because he anticipates and discusses both the inward and outward impacts in the long-term future of PP as a whole, as well as the possible follow up projects and the clinical implementation of the results of research in precision medicine. He is able to do so because he possesses a deep knowledge not only of the biomedical research field, but also of the political and institutional context in which research is embedded. Matthew describes his own work as the struggle to constantly “think big, raise up, look further toward the horizon and make plans”.

Another person, ‘Luke’, was able not only to discuss the project pipeline as a whole, but also to look at the whole field of precision medicine from a wider perspective and to critically discuss it in the context of current and even future biomedical research. Since he puts the whole PP into a wider perspective, thus surpassing even the forward-looking researchers in terms of anticipative reflection, I defined him as possessing an *eagle-eye view*. For Luke, in the medium to long-term future, other approaches to oncology will be more fruitful than precision medicine. At the same time, he is also able to provide a convincing pragmatic argument for the need of continuing to do research in precision medicine now.

I: *Why have you chosen to work in a project about precision medicine if you think there is a better alternative available?*

Luke (PI): “The reason for that is that these immunotherapy approaches, they will be better, but it will take ten, or twenty, or maybe even thirty years before they are fully operational. Until then, we have to make tools with what we have and we have all these drugs that are approved for various cancers and they probably work for certain patients, but then we

need to find out who are these patients, who benefits from these compounds, who benefits from combinations of compounds and who does not. So in the medium long term I think precision medicine can fill the gap until we have developed better tools for cancer treatments”.

An actual scientist as reflective as Luke does not fit neatly with the abstract idea of a socially responsible scientist. While he is indeed capable of engaging with future-oriented reasoning, and even more so than Matthew, he does not regard being engaged with the public and other stakeholders as one of his duties, nor he appears to be particularly driven by moral and social concern, apart from being a ‘good researcher’ (which involves the production of reliable knowledge). Under this respect, Luke’s attitude seems specular to Mark’s, who is driven by moral and social values but who does not engage in anticipative reflection (and, moreover, thinks that it is responsible not to engage in it).

Thomas, on the one hand, and Mark, Matthew, and Luke, on the other hand, occupy the opposite extremes of the spectrum of the degrees of reflection about the future implications of their research. They are also at the two extremes of the academic hierarchy: as already said, while Thomas is a junior scientist without a permanent position, the other three are tenured PIs.

6.4 We-reasoning

The results analysed so far concern individuals within PP: whether they are able to anticipate the future impacts of their research, and of which kind, and whether some of them deviated from the others in their engagement with anticipative thinking. The answers of nine PP researchers, however, signal the presence of the form of ‘we-reasoning’ that, as explained in section 2, is necessary to collective agency.

Two specifications are in order. First, the *We-reasoning*-theme mostly emerged in relation to questions about methodological choices: how they are made and who makes them. This means that PP-researchers consider themselves involved in the ‘we-reasoning’ mainly when it comes to make decisions aimed at reaching their *epistemic* goals. Second, by going more in depth into these issues and by asking them to clarify their answers, interviewees were able to distinguish how different individuals contributed to the collective decision-making process. The collective ‘we’ is not an

abstract entity that goes above and beyond the members of a sub-team, and the decisions it takes are the outcome of deliberative processes in which different individuals contribute differently.

Julie (mid-level researcher): “*We analyse things and, if there are mistakes, then we need to repeat, or maybe if it’s not possible to repeat we need to find... we need to analyse a bit better to, kind of, say... we need to trust this, or maybe it can be used in another way?, or... so... depends. There are many different types of interpretations of your results, I think*”.

I: *And who decides.? Who does the interpretation work? Who interprets the results?*

Julie (mid-level researcher): “*That is primarily done by the person who is doing the experiment. In this type of project you really work together with the biostatisticians, because I don’t process the raw data coming directly from the machine, so you get numbers that I need to reinterpret and then again I do my interpretation that I show in the group meeting, to Matthew and other researchers, so that the interpretation is done together with him. So, yeah. I’m not taking, you know, the responsibility alone to publish something, but I discuss*”.

‘Charlotte’, a senior researcher, was able to distinguish the different roles that she and the other members of her sub-team play in making decisions, based on their level of competence and seniority. Even if she makes some of the ‘final’ decisions, the ways in which she arrives at such decisions is heavily informed by discussions with all the other individuals working with her.

I: *You and your group are developing the software and also thinking about the underlying mathematical approach [...]. Try to explain to me, as someone from the outside, how would you choose the right mathematical approach? How many alternatives do you have? And how difficult is it to make such a choice?*

Charlotte (senior researcher): “That is a good question!”

I: *I often have this reaction in these interviews. Thank you.*

Charlotte (senior researcher): “Do you? Well, this is not an obvious question. So, there are some obvious types of models which lend themselves to this task, and that *we* have more than others, and maybe [*we*] make the decision that *we* use one of them, and *we* try to write down the specifications that the model needs to fulfil, right? One thing that *we* know, for example, is that a lot of the relationships are highly non-linear, they are kind of complex relations, so the model needs to be able to deal with that. Then *we* need to be able to scale to a large dimension, kind of a bigger scene, so it’s also a computational issue, because *we* cannot go... *we* are forced to look at the higher level, because *we* look at the big screens so it’s computational and feasible, otherwise. And then there is, yes, *we* just follow the literature and actually try to follow the literature in different fields, which is also challenging. So, the machine learning field, which is just coming from the computer science side, then also the statistical field, and these recent developments in both fields... and these communities overlap, but they publish in different journals, so *we* have to follow them two. Yeah. And then try to kind of make an ‘educated guess’, formulate a project based on that. *We* are also updating and learning over time, right?, but it’s also...”

I: *You said: “We try to make an educated guess”, but who is ‘we’? Who is making the guess?*

Charlotte (senior researcher): “[*laughing*] *It depends.* When I write the project description for a PhD student, I would probably usually decide what *we* are going to look at. Then there will be a refinement in the project, but the main approaches are given at the beginning. But for a postdoc like F., for example, I have been keeping that quite open, *we* discuss things together. I make suggestions based on the experiences *we’ve* had with Sten’s project, and A.’s projects, and also another PhD student who was working in another related area, and so on with all the experiences *we’ve* had... But actually, *in the end, probably a lot, it’s a lot my decision*”.

When it comes to the attribution of responsibility, however, the collective nature of research may pose some dilemmas. For example, ‘Sten’ seemed to maintain that blame for something going wrong should go to someone in particular, whereas praise for something good should go to the whole team.

I: *In these long projects there are always people coming and going. In your view, if anything goes wrong, who should be held responsible for the mistake?*

Sten (PhD student): “I don't know. It's a difficult question. Depends on what you mean by ‘going wrong’, so... *We* do as much as *we* can to test the software to make sure that it's functioning in the reasonable data cases, and *we* provide help files, and vignettes, and *we* write papers to explain how the package works and how to fix errors. *We* use a very particular probabilistic language, which is the software the package is built on, which is heavily developed and supported, it has an active community. Everything is open source, so you can see all the files, how the model is set up and everything. In the end who can be held responsible... It depends on what goes wrong”.

I: *You told me about things going wrong, about the tool and the biostatistical models, but something can go wrong with the project, which involves different people. If anything in general, with the project, goes wrong do the biostatisticians say: ‘Ah, we only do numbers, we develop the models, it's not our responsibility’.*

Sten (PhD student): “I don't know. It's a package deal. *We* are all trying to do the best *we* can”.

I: *Sure, mine is not a comment about the intentions. Of course, the intentions are always very good, and the efforts as well.*

Sten (PhD student): “I am sure there is probably a correct answer here, because there is a pot of money that's given to *someone* under some agreement, under some deliverables. If we don't do anything, we just sit back and get paid and don't do any experiment, don't do anything, *someone* will be held responsible for that.

I: *Let me put the question from another side: if everything goes perfectly right, who is going to be responsible for the success of the project? One or two people, or the whole team?*

Sten (PhD student): “[*thinking about it*] The whole team, but maybe in different ways”.

When he was made aware of this asymmetry, ‘Sten’ added that responsibility should be distributed in different ways across the group, but he was not sure of how such a distribution should proceed.

I: *So, in the case of success everybody is good? I’m not even sure there is a right answer, there’s not even this assumption here. But so, if everything goes ok, everybody is good and everybody’s is a success story, but if anything goes wrong then the responsibility is just of someone... Then you said there’s a pot of money given to do the research, but part of that money went to your pocket as well, because you are also at the end line of the money.*

Sten (PhD student): “Definitely yes”.

I: *So, it’s not like a one person thing.*

Sten (PhD student): “*I think we are all involved, but we have different levels of responsibility. But of course if you do your job and the project is successful because of my little piece of input, then I’m happy of course. If things go wrong and it’s maybe my fault for building a tool that is terrible and doesn't work, I would be sad of course*”.

The discussion with ‘Sten’ mainly revolved around accountability, that is about backward-looking responsibility. However, the same difficulties in deciding who is supposed to be the responsible agent of PP, whether the group as a whole or some individuals in particular, also arose in discussions about forward-looking responsibility. ‘Matthew’, the forward-looking PI of section 6.3, would like a major involvement of every member of PP in the kind of anticipative reflection he is already conducting. He also hopes to encourage such a reflective attitude by leading through examples. Despite his optimism, when asked whether the other members of PP are actually starting to ‘look towards the horizon’, ‘Matthew’ realises that anticipative reflection is not easy to be taught by example.

Matthew (PI): “I do think the default way that I would like to lead people in research is by ‘charismatic leadership’. You sell a vision to the other members of the project and you set up the goal a little bit further up in the horizon as something which would be a good place, a goal. And if it’s difficult to get there you also start groundbreaking projects, you know. [...] You sell a vision and you get other people to buy into your vision, because that involves setting up something towards the horizon”.

I: *It seems to me that, as a PI, you have a very wide outlook on things, you think forward and very much so. You are already thinking about what will come next, the main problems with implementation, who needs to integrate what. You have these problems very clearly in mind. Do you think that [the other members of your team] have also got this wide view on research, or is their view a bit more on the ‘here-and-now’?*

Matthew (PI): “[with hesitation] I think they are a bit more on the here-and-now”.

7. Discussion: Distributing responsibilities

PP members appear far from the ideal (or stereotype) of the ‘disinterested scientist’. As discussed in 6.1, many of them are driven by a desire to improve society and to have a positive impact on cancer patients. However, intention is only one of the necessary conditions for the attribution of responsibility. One of the most striking results of this study is the evident gap between researchers’ value-drivenness and good intentions, on the one hand, and their little propensity to think about the potential outward impacts of their work, or to see themselves as the agents operating such kinds of transformative impacts.

This does not mean that no one, at PP, is engaging with the kind of anticipative reflection that is considered necessary to the social responsibility of science. This is the case of some of the individuals displaying deviations from the rest of the group. In short, while the majority of PP members, especially junior and mid-level researchers, carry on the epistemic responsibility of producing justified and reliable information, and therefore of thinking about the inward impacts of the project, the more experienced PIs seem to carry the duty of ‘thinking about the future’ and reflecting about the wider social and institutional context of their research.

The qualitative data collected during the study, however, resist the over-simplistic explanation based on power dynamics within the academic hierarchy. On the one hand, some parts of the PP pipeline (that is, some of its sub-teams) are actually characterised by a rather ‘horizontal’ structure in which research is conducted in a strongly collaborative and mutually respectful spirit. ‘Mark’, the questioning researcher, actively encourages the members of his team to be reflective about the value and impact of research. ‘Matthew’, the forward-looking researcher, would prefer social responsibility to be equally shared among all the members of PP and, more than a decision-maker, he regards himself as someone who leads by example. Yet, at the same time, he recognizes the difficulties of inculcating a sense of anticipative reasoning in his collaborators.

On the other hand, many PP-researchers do have real concerns about social and moral issues. When explicitly asked why they do not think more, and more deeply, about the social implications of their work, *despite* their own motivations and good intentions, some of them replied that they were not ‘qualified’ or ‘prepared enough’ for undertaking such a task. They felt that their PIs were far better suited to think about ‘big issues’. What seems to regulate the distribution of responsibilities in PP, in short, is the *feeling of inadequacy* perceived by many researchers. It is for this reason that they prefer to delegate the engagement with the reflection on the moral and social implications of their research to the more experienced PIs.

This is not a problem per se because, following some theories on collective responsibility, for a group to be collectively responsible it is not necessary that every single one of its members carries the same duties in the same way at the same time. Theories such as those developed by List and Pettit (2011) or Collins (2019) allow the assignment of special duties to those individuals who, within a group, are the most capable of carrying them. Such individuals, in turn, represent the group they are part of, meaning that their decisions ought to reflect its collective ‘ethos’, which also unites the other individual members.

In the context of philosophy of science, Rolin (2017) has borrowed some concepts from moral philosophy to develop what could be defined as a distributive approach to the issue of scientists’ responsibilities. Since it would be unjust, if not unrealistic, to expect that every single scientist is capable of carrying different kinds of responsibilities in the same way and at the same time, Rolin suggests that only few members of the community ought to carry special duties. She explains that, while *general responsibilities* are universal, and everybody should carry them, so-called *special*

responsibilities are general responsibilities assigned to individuals who can play a specific role within their community. To make an example: while, as a general rule, everybody has the duty to try to rescue people in danger, in some circumstances (i.e., at a beach) only few qualified people (i.e., lifeguards) have the special responsibility of carrying such a duty. Similarly, for Rolin, it is possible to conceive a scientific community in which only few individuals carry some of the community's duties.

Following on these insights, PP seems to display a *spontaneous* and *implicit* distribution of responsibilities. While divisions of practical tasks and research duties occur regularly in research groups, and are actually part of the planning of the project, the distribution of the 'moral work' at PP is implicit in the sense that it is not the outcome of planning or deliberation. The main issue, however, is not that responsibilities are distributed within PP, but whether they are distributed in the right way. An implicit distribution of responsibilities, in fact, could be problematic for a number of reasons.

To begin with, while it is true that senior researchers may be more aware of the social and institutional context in which science is embedded, and may have also acquired the right skills, or even the authority, to engage responsibly with different stakeholders, knowledge and awareness are not in themselves sufficient for responsibility. People can learn new things, and some people may know more things than others, yet knowing more does not automatically change how one thinks or acts. Some people may even use what they know for pursuing selfish or even immoral agendas. Because knowledge is only one of the necessary conditions for responsibility, irresponsibility is not necessarily and not always caused by a deficit in knowledge. In short, assuming that some individuals occupying a high position in a hierarchy 'know better' only because they 'know more', and delegating to them important responsibilities on the basis of such an assumption, is unwarranted and problematic.

Nor should it be assumed that younger researchers do not possess any valuable knowledge or a future-oriented attitude. The anticipative reflections of these individuals, however, will have little or no bearing in the directions taken by research, especially if they will feel like their own contribution is not relevant enough. In moral philosophy, this is known as the *individual impotence objection*, for which individuals feel dispensed from partaking in the performance of collective duties (Parfit 1984; Pinkert 2015). Perceived individual impotence clearly arises at the *collective* level and it

is a problem that the interventions at the midstream stages seeking to encourage anticipative reflection among *individual* scientists do not consider enough. The issue is not whether individuals engage with a particular type of moral reasoning or action, but whether they think that their reasoning or action may be impactful in the grand scheme of group dynamics.

Moreover, a just distributive model of responsibilities should not correspond to the uncritical delegation of special duties to few members of the group, completely dispelling the others from any participation in the performance of such duties. Doing so would reinforce the diffuse perception of personal impotence, while depriving the group of potentially valuable inputs from individuals who do *not* have special roles or special duties. To put the whole weight of a special responsibility on the shoulders of only a few individuals would hardly contribute to making the group as a whole ‘collectively responsible’.

These issues may be reinforced by the institutional elements of scientific research. On the one hand, PIs have to write research proposals and look for funding. They are often required to write ‘impact statements’, which forces them in some respects to engage with anticipative reflection. On the other hand, junior researchers rarely get credit for thinking about these issues: for them, epistemic success counts primarily. Even in cases such as PP, characterised by ‘horizontal’ research structures, the institutional dimension governing research may end up creating hierarchical splits in matters related to social responsibility.

Some of the interventions aimed at changing the institutional dimension of science do not do much either. For example, Digital Life Norway, which funded PP, organises RRI intensive courses for graduate students and junior researchers. These training activities clearly target only part of the population of scientists. While many of the people who took the RRI courses will leave academic research, or even research altogether, the implicit distribution of responsibilities, with junior researchers doing the benchwork and PIs thinking about future impacts, remains in place *de facto*.

These observations lead to the final point of this discussion, about whether PP is a collective agent, and of which kind. PP appears like a collective *epistemic* subject, characterised by a set of shared epistemic goals, and of collaborative and coordinated activities to reach them. More precisely, PP appears like a collection of small collective epistemic subjects (the sub-teams working within the pipeline), characterised by the

‘we-reasoning’ and in which decisions are taken at the collective level; the sub-teams, in turn, form a bigger collective epistemic subject (the whole working pipeline). When it comes to moral agency, however, PP could be at best defined as a ‘semi-collective’ agent, in virtue of the fact, discussed above, that the distribution of special roles is implicit and perhaps not entirely just, or even effective.

Whether such a (semi)collective agent is of the kind theorised by Bratman or it is, rather, closer to the kind discussed by philosophers like Collins or List and Pettit should be the object of further analysis. What is important to stress is that while PP could be further ‘collectivised’ through a more explicit and transparent distribution of special duties, this would not automatically imply that its members will get more ‘responsibilized’. The ‘we-reasoning’, characteristic of collective agents, may be used by individuals to de-responsabilise themselves, by avoiding to regard themselves as responsible for some of the decisions that are taken at a collective level. This means that the interest of the individual members of a group towards their collective responsibilities should not stop once such responsibilities have been distributed in an explicit and just way.

8. Doing Social Epistemology of Responsible Science *via* qualitative studies

The issue of the social responsibility of science is complex and multi-faceted. The problem of which kind of group structure best serves the kind of collective remedial responsibility that current science policies demand from research groups requires the development of more detailed normative theories. The aim of this article was not to develop such a normative theory in full. Rather, I used a qualitative case study to show what the current debate about the social responsibility of science lacks. Such a debate is often framed around the concepts of ‘the responsible science’ or of ‘the responsible community’, both treated in rather idealised terms. These perspectives however fail to capture the complexity of group dynamics in actual research groups. The case study has been used as an empirical ground to highlight the set of problems that the collective dimension of scientific research poses to the emerging conception of a socially responsible science, and that the philosophical debates should attempt to solve. Even from the description of the case study, moreover, it is possible to begin to draw some prescriptions. For example, the analysis of an implicit distribution of duties within a research group invites us to reflect on *how* such a distribution of duties ought to be properly operated.

This approach is not a novelty. Several philosophers have already used qualitative research methods for philosophical purposes (see, for example, MacLeod 2018; MacLeod and Nersessian 2013, 2015, 2016; Nersessian 2019, 2022; Nersessian et al 2003a,b; Osbeck and Nersessian 2017; Wagenknecht 2016; Hangel and Schickore 2017; Schickore and Hangel 2019; Trappes 2022). The motivation for employing such methods is that they may represent the missing link between general prescriptions and the details of actual scientific practice, and may therefore be extremely useful for philosophical research (Nersessian and MacLeod 2022; Hangel and ChoGlueck 2023).

This article contributes to the qualitative research approach in philosophy of science by presenting a case study about the distribution of social responsibility in a research team, representative of the kind of research that is funded by agencies that require scientists to be socially responsible and to comply with RRI. While they acknowledge the role of subjective factors, such as researchers' own standpoints, many of the qualitative studies in philosophy of science conducted so far have focused on the *epistemic* elements of scientific research (i.e., how scientists reason to construct a model, how they define or operationalise specific concepts, and so on). In this way, these studies remain within the same scope of traditional social epistemology of science, which seems to miss a 'moral dimension'.

At the same time, and as noticed by de Melo-Martin and Intemann (2023), many philosophical debates about values in science and responsibility are conducted from a rather individualistic perspective. Such debates, in fact, attempt to develop normative frameworks for the (ideal) individual scientist, regarded as the fundamental responsible agent in science. However, group dynamics may constrain individual responsibilities in different ways, the distribution of responsibilities discussed in this article being an example. In short, the debate about the responsibility of science seems to miss a 'social dimension'.

The present study complements both empirical social epistemology and the philosophical debate on the responsibility of science, laying the ground for what I define as the *social epistemology of responsible science*.

There are, of course, potential issues with the use of a qualitative case study for drawing normative claims. One of the perceived issues of qualitative studies is that, focused as they are on 'local' analyses, and unlike quantitative research, their results are not statistically valid and cannot be generalised. It is therefore unclear how context-specific qualitative studies may be used to support general philosophical

conclusions. Some philosophers respond by noticing that this problem is akin to the one posed by the use of historical case studies in philosophy of science. That historical case studies are not generalizable strictly speaking does not prevent philosophers from using them to develop general frameworks. Following the view of Chang (2012) on the relation between particular episodes in the history of science and general philosophical claims, Masneures and Wagenknecht (2015) argue that qualitative case studies and philosophical theories are not in a “particular-to-general” relation but, rather, in a “concrete-to-abstract” relation.

Such a consideration complements the position elaborated by many theorists of sociological and qualitative research. Already Weber (1949) argued that the object of sociological explanations are ‘ideal types’ that, even though are never fully instantiated in reality, are ‘grounded’ on empirical data (which, in turn, allow the researcher to find and refine the ideal categories). Similarly, the so-called ‘sensitising concepts’ used in contemporary qualitative research are developed by abstracting from concrete data and are used as heuristic devices: they “provide a place to *start* inquiry, not to *end* it” (Chamaz 2014:31), and can be tested and further developed in future research (Blumer 1954; Bowen 2006). More generally, supporters of the use of qualitative case studies also argue that, even though they produce context-dependent information, they represent concrete ‘exemplars’ that other researchers may use for addressing some problems in similar contexts (Flyvbjerg 2004). This is why social theorists argue that the most important trait of case studies is not their generalizability, but the validity of the analysis they propose, which in turn allows logical inferences, which are different from statistical inferences (Mitchell 1983).

In short: qualitative studies in philosophy of science are a mainly *descriptive* enterprise of actual local scientific knowledge production; such descriptions, however, may also play a *heuristic role* (in the sense that they may lead to the discovery of new philosophically interesting facts about scientific practice) and even a *normative role* (in the sense that they may challenge some of the existing normative theories and lead to the examination of some norms that are ‘implicit’ in the actual practice). In the present article, I have discussed a qualitative study with the purpose of laying the ground for a social epistemology of responsible research.

9. Conclusions

Remedial responsibility requires the ability to imagine a desirable future state of affairs and anticipate the future impacts of the present course of actions. What exactly anticipative reflection is and how people engage with it are issues that would enormously benefit from the contributions of further studies. Within the limited scopes of a single qualitative case study, in this article I have focussed on how anticipative reflection is carried on within an actual research group. The results show a more or less implicit distribution of duties and roles, which however may not be sufficient for ascribing collective moral responsibility to the group as a whole. For this reason, it is necessary to complement existing social epistemological analyses of the social structure of scientific research with further studies on the collective nature of the social responsibility of science.

Acknowledgements: I am grateful to two anonymous reviewers, whose insightful comments helped me improve this article considerably. I will always be grateful to the members of the PINpOINT project for the invaluable experience and for welcoming a philosopher with a collaborative spirit and open mind. This research was funded by the Centre for Digital Life Norway (Grant number: 294916). It was completed thanks to the support of the Beatriu de Pinós/Marie Skłodowska–Curie Action Cofund Fellowship (Grant number: 2020-BP-00196).

References

- Barbolosi, D., Ciccolini, J., Lacarelle, B., Barlési, F., & André, N. (2016). Computational oncology: Mathematical modelling of drug regimens for precision medicine. *Nature Review Clinical Oncology*, 13, 242–254.
- Berenbaum, M. (1989). What is synergy? *Pharmacological Reviews*, 41, 93–14.
- Blumer, H. (1954). What is wrong with social theory? *American Sociological Review*, 18, 3–10.
- Bowen, G. (2006). Grounded theory and sensitising concepts. *International Journal of Qualitative Methods*, 5, 12–23.
- Bratman, M. (1992). Shared cooperative activity. *The Philosophical Review*, 101, 327–341.
- Bratman, M. (1993). Shared intentions. *Ethics*, 104, 97–113.

- Bratman, M. (2013). *Shared agency: A planning theory of acting together*. Oxford University Press.
- Carrier, M. (2021). How to conceive of science for the benefit of society: Prospects of responsible research and innovation. *Synthese*, 198, 4749–4768.
- Chang, H. (2012). Beyond case studies: History as philosophy. In S. Mauskopf & T. Schmaltz (Eds.), *Integrating History and Philosophy of Science: Problems and Prospects* (pp. 109–124). Routledge.
- Charmaz, K. (2014). *Constructing grounded theory*. Sage.
- Collins, S. (2019). *Group duties: Their existence and implications for individuals*. Oxford University Press.
- Corlett, J. (2001). Collective moral philosophy. *Journal of Social Philosophy*, 32, 573–584.
- de Melo-Martín, I., & Intemann, K. (2023). Socially responsible science: Exploring the complexities. *European Journal for Philosophy of Science*, 13, 33. <https://doi.org/10.1007/s13194-023-00537-6>
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. Pittsburgh University Press.
- European Commission. (2014). *Science with and for Society*. Available at: <https://ec.europa.eu/programmes/horizon2020/en/h2020-section/science-and-society>
- Gilbert, M. (1990). Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy*, 15, 1–14.
- Gilbert, M. (2000). *Sociality and responsibility*. Rowman and Littlefield.
- Gilbert, M. (2013). *Joint commitment*. Oxford University Press.
- Goldman, A. (1970). *A theory of human action*. Princeton University Press.
- Figueiredo Nascimento, S., Cuccillato, E., Schade, S., & Guimarães Pereira, A. (2014). *Citizen Engagement in science and policy-making*. Publications Office of the European Union.
- Fisher, E., Mahajan, R., & Mitcham, C. (2006). Midstream modulation of technology: Governance from within. *Bulletin of Science, Technology and Society*, 26, 485–496.
- Fisher, E. & Rip, A. (2013). Responsible Innovation: multi-level dynamics and soft intervention practices. In Owen R., Bessant J. & Heintz M. (Eds). *Responsible*

- Innovation. Managing the Responsible Emergence of Science and Innovation in Society*. Chichester: Wiley, 165–183.
- Flyvbjerg, B. (2004). Five misunderstandings about case-study research. *Qualitative Inquiry*, 12, 219–245.
- French, P. (1984). *Collective and corporate responsibility*. Columbia University Press.
- Funtowicz, S., & Ravetz, J. (1993). Science for the post-normal age. *Futures*, 31, 735–755.
- Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P., & Trow, M. (1994). *The New Production of Knowledge: The dynamics of science and research in contemporary societies*. SAGE.
- Greco, W., Bravo, G., & Parsons, J. (1995). The search for synergy - a critical review from a response-surface perspective. *Pharmacological Reviews*, 47, 331–385.
- Hangel, N., & ChoGlueck, C. (2023). On the pursuitworthiness of qualitative methods in empirical philosophy of science. *Studies in History and Philosophy of Science*, 98, 29–39.
- Hangel, N., & Schickore, J. (2017). Scientists' conceptions of good research practice. *Perspectives on Science*, 25, 766–791.
- Haraway, D. (2016). *Staying with the Trouble: Making Kin in the Chthulucene*. Duke University Press.
- Intemann, K. (2015). Distinguishing between legitimate and illegitimate values in climate modelling. *European Journal for Philosophy of Science*, 5, 217–232.
- Intemann, K., & de Melo-Martin, I. (2010). Social values as evidentiary standards: The case of the HPV vaccine. *Biology and Philosophy*, 25, 203–213.
- Jasanoff, S. (2016). *The Ethics of Invention: technology and the human future*. W.W. Norton & Company.
- Kitcher, P. (1990). The Division of Cognitive Labor. *Journal of Philosophy*, 87, 5–22.
- Kitcher, P. (2001). *Science, Truth, and Democracy*. Oxford University Press.
- Kitcher, P. (2011). *Science in a Democratic Society*. Prometheus Books.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *The Philosophical Review*, 92, 33–48.
- Kourany, J. (2010). *Philosophy of Science after Feminism*. Oxford University Press.
- Lewis, H. (1948). Collective responsibility. *Philosophy*, 24, 3–18.
- List, C., & Pettit, P. (2011). *Group Agency: The possibility, design, and status of corporate agents*. Oxford University Press.

- Longino, H. (1990). *Science as social knowledge*. Princeton University Press.
- Longino, H. (2002). *The fate of knowledge*. Princeton University Press.
- Masneures, E., & Wagenknecht, S. (2015). Feeling with the organism: A blueprint for empirical philosophy of science. In S. Wagenknecht, N. Nersessian, & H. Andersen (Eds.), *Empirical Philosophy of Science: Introducing Qualitative Methods into Philosophy of Science* (pp. 37–64). Springer.
- Mitchell, J. (1983). Case and situation analysis. *The Sociological Review*, 31, 187–211.
- MacLeod, M. (2018). What makes interdisciplinarity difficult? *Synthese*, 195, 697–720.
- MacLeod, M., & Nersessian, N. (2013). Building Simulations from the Ground Up: Modelling and Theory in Systems Biology. *Philosophy of Science*, 80, 533–556.
- MacLeod, M., & Nersessian, N. (2015). Modelling systems-level dynamics: Understanding without mechanistic explanation in integrative systems biology. *Studies in History and Philosophy of Science c: Studies in History and Philosophy of Biological and Biomedical Sciences*, 49, 1–11.
- MacLeod, M., & Nersessian, N. (2016). Interdisciplinary problem-solving: Emerging modes in integrative systems biology. *European Journal for Philosophy of Science*, 6, 401–418.
- Nersessian, N. (2019). Interdisciplinarity in action: Cognitive ethnography of bioengineering sciences research laboratories. *Perspectives on Science*, 27, 553–581.
- Nersessian, N. (2022). *Interdisciplinarity in the Making*. MIT Press.
- Nersessian, N., Kurz-Milcke, E., Newstetter, W., & Davies, J. (2003a). Research laboratories as evolving distributed cognitive systems. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, 857–862.
- Nersessian, N., Newstetter, W., Kurz-Milcke, E., & Davies, J. (2003b). A mixed-method approach to studying distributed cognition in evolving environments. In *Proceedings of the International Conference on Learning Sciences*, 307–314.
- Nersessian, N., & MacLeod, M. (2022). Rethinking ethnography for philosophy of science. *Philosophy of Science*, 89, 721–741.
- Nordmann, A. (2014). Responsible Innovation, the art and craft of anticipation. *Journal of Responsible Innovation*, 1, 87–98.

- Osbeck, L., & Nersessian, N. (2017). Epistemic identities in interdisciplinary science. *Perspectives on Science*, 25, 226–260.
- Owen, R., Stilgoe, J., Macnaghten, P., Gorman, M., Fisher, E., & Guston, D. (2013). A framework for Responsible Innovation. In R. Owen, J. Bessant, & M. Heintz (Eds.), *Responsible Innovation* (pp. 27–50). John Wiley.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Pinkert, F. (2015). What if I cannot make a difference (and I know it). *Ethics*, 125, 971–998.
- Ratti, E., & Russo, F. (2024). Science and values: A two-way direction. *European Journal for Philosophy of Science*, 14, 1–23.
- Rolin, K. (2017a). Scientific community: A moral dimension. *Social Epistemology*, 31, 468–483.
- Rolin, K. (2017b). Scientific dissent and a fair distribution of epistemic responsibility. *Public Affairs Quarterly*, 31, 209–228.
- Rovane, C. (1997). *Bounds of agency: An essay in revisionary metaphysics*. Princeton University Press.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20, 1–6.
- Schickore, J., & Hangel, N. (2019). “It might be this, it should be that...” uncertainty and doubt in day-to-day research practice. *European Journal for Philosophy of Science*, 9, 1–21.
- Schwenkenbecher, A. (2019). Collective moral obligations: ‘we-reasoning’ and the perspective of the deliberative agent. *The Monist*, 102, 151–171.
- Selin, C. (2014). On not forgetting futures. *Journal of Responsible Innovation*, 1, 103–108.
- Smiley, M. (2022). Collective responsibility. In Zalta, E. & Nodelman, U. (Ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2023 edition), URL=<https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=collective-responsibility>
- Sverdlik, S. (1987). Collective responsibility. *Philosophical Studies*, 51, 61–76.
- Timmermans, J., & Blok, V. (2021). A critical hermeneutic reflection on the paradigm-level assumptions underlying responsible innovation. *Synthese*, 198, 4635–4666.

- Trappes, R. (2022). Individual differences, uniqueness, and individuality in behavioural ecology. *Studies in History and Philosophy of Science*, 96, 18–26.
- Urueña, S. (2021). Responsibility through Anticipation? The ‘Future Talk’ and the Quest for Plausibility in the Governance of Emerging Technologies. *NanoEthics*, 15, 271–302.
- van der Burg, S. (2009). Taking the ‘soft impacts’ of technology into account: Broadening the discourse in research practice. *Social Epistemology*, 23, 301–316.
- van der Burg, S. (2014). On the hermeneutic need for future anticipation. *Journal of Responsible Innovation*, 1, 99–102.
- Von Schomberg, R. (2013). A vision of responsible innovation. In: R. Owen, M. Heintz and J Bessant (Eds.) *Responsible Innovation*. London: John Wiley.
- Wagenknecht, S. (2016). *A Social Epistemology of Research Groups*. Palgrave MacMillian.
- Watkins, J. (1957). Methodological individualism and social tendencies. *British Journal for the Philosophy of Science*, 8, 104–117.
- Weber, M. (1949). *The Methodology of the Social Sciences*. E. Shils and H. Finch (Eds.), Free Press.
- Wildson, J. (2014). From foresight to hindsight: The promise of history in responsible innovation. *Journal of Responsible Innovation*, 1, 109–112.
- Wilholt, T. (2016). Collaborative research, scientific communities, and the social diffusion of trustworthiness. In M. Brady & M. Fricker (Eds.), *The Epistemic Life of Groups: Essays in the Epistemology of Collectives* (pp. 218–234). Oxford University Press.
- Zahle, J. (2023). Reactivity and good data in qualitative data collection. *European Journal for Philosophy of Science*, 13, 10.
<https://doi.org/10.1007/s13194-023-00514-z>

	Future	Mode	Responsibility	Based on	Agent
Inward Impact	Short/Medium	Definite	Epistemic	Knowledge, Expertise	'We'
Outward Impact	Long	Vague	Social/ Remedial	Hopes, Intuitions	Others

Table 1: Inward and outward impacts anticipated by the members of PP.