

What is Collective Responsibility?

Rowan Mellor

University College London

Philosophical Studies MPhil Stud

Declaration

I, Rowan Mellor, confirm that the work presented in this thesis is my own.
Where information has been derived from other sources, I confirm that this
has been indicated in the thesis.

Abstract

This thesis attempts to answer the following question; does an understanding of collective responsibility require different conceptual resources to those employed in an account of singular responsibility? I argue that the answer to this question can be 'no'. I do so by assuming a particular theory of singular responsibility. I then argue that collective responsibility, in its various forms, is consistent with this theory. The thesis is divided into four chapters. Chapter 1 presents the central problem of the thesis and gives a preliminary discussion of responsibility and plurals. In Chapter 2 I discuss two notions of collective responsibility; collective responsibility as several responsibility and collective responsibility as corporate responsibility. Chapter 3 discusses a third possible notion of collective responsibility; collective responsibility as plural responsibility. In Chapter 4 I present some cases in which collective responsibility as plural responsibility might occur.

Acknowledgements

Thanks to everyone who helped with this project and to those who read the endless papers that may, or may not, have ended up here in some form.

Special thanks go to my primary supervisor Véronique Munoz-Dardé. Thanks also to Doug Lavin, Mike Martin and Han van Wietmarschen for lengthy discussions on these topics. Finally thanks to those who attended the 2016 Brave New World Conference at the Manchester Centre for Political Theory and the Moral and Political Philosophy Workshop at UCL. Draft chapters of this thesis were presented at both sessions.

Contents

Introduction	7
1. Responsibility and Plurality	9
1.1. Responsibility	9
1.1.1. Substantive responsibility	10
1.1.2. Holding responsible	11
1.2. Responsibility as self-governance	13
1.2.1. Reasons	14
1.2.2. Self-governance	16
1.3. The problem of collective responsibility	18
1.4. Plurality	22
1.4.1. Plural logics	22
1.4.2. Property possession	24
1.4.3. Action	27
1.5. Conclusion	28
2. Several and Corporate Responsibility	30
2.1. Several responsibility	30
2.1.1. Lewis' challenge	31
2.1.2. Responses to Lewis	33
2.2. Corporate responsibility	39
2.2.1. Collective agents	40
2.2.2. Corporate and several responsibility	43
2.2.3. Corporate responsibility and non-agents	45
2.3. Conclusion	48
3. Plural Responsibility	49
3.1. Plural responsibility and self-governance	49
3.1.1. Agents and actions	50
3.1.2. Plural self-governance	52
3.2. Plural reasons	54
3.2.1. What is a plural reason?	55
3.2.2. Plural reasons and self-governance	57

3.3. Reasoning about plurals	59
3.3.1. Sellars' Problem.....	60
3.3.2. Some solutions.....	61
3.4. Conclusion.....	65
4. Cases of Plural Responsibility	67
4.1. Coordination problems.....	67
4.2. Jackson's case.....	72
4.3. No difference cases.....	75
4.3.1. No difference cases and singular reasons.....	77
4.3.2. No difference cases and plural reasons	83
4.4. Consequentialism.....	85
4.5. Conclusion.....	86
Conclusion.....	88
Bibliography	90

Introduction

We are familiar with the idea of someone being responsible for some event. Additionally, I think we have become familiar with the possibility of several people being *collectively responsible* for some event(s). But what is collective responsibility? Under what conditions does it occur? Do we need to expand our ordinary notion of responsibility in order to understand collective responsibility? Does a commitment to collective responsibility presuppose the truth of certain ontological claims? The aim of this thesis is to go some way towards answering such questions. The primary argument will be that our ordinary notion of responsibility need not be revised in order to accommodate collective responsibility. Moreover a commitment to certain forms of collective responsibility presupposes minimal ontological commitments.

Collective responsibility is most frequently invoked in political contexts. Consequently, many discussions of collective responsibility have overt political ambitions.¹ This thesis is not a work of political theory. Rather, I view it as a study in moral philosophy. My aim is to analyse the moral concept ‘collective responsibility’; to spell out some of its forms, to trace its connections to other areas of moral life and to trace any theoretical commitments it might presuppose in other areas of philosophy. However my project does have political foundations. The questions with which this thesis is concerned arose in response to Iris Marion Young’s (2011) posthumously published work on structural injustice. Young suggests that, by virtue of participating in social-structural processes, millions of individuals are causally implicated in the perpetuation of various injustices such as sweat-shop labour, homelessness and gender inequality. In order to theorise about responsibility in relation to such injustices, Young argues that we need a model of responsibility on which “[m]y responsibility is essentially shared with others” (Young 2011, p. 110). Young and I differ in our treatments of responsibility (Young seems primarily concerned with what I refer to as ‘substantive

¹ See e.g. (Arendt 1987; Jaspers 2000; May 1992).

responsibility'²). However the central arguments of this thesis are an attempt address a question which Young neglects; what is it for responsibility to be shared?

The thesis is divided into four chapters. Chapter 1 presents the central problem of the thesis and gives a preliminary discussion of responsibility and plurals. In Chapter 2 I discuss two notions of collective responsibility; collective responsibility as several responsibility and collective responsibility as corporate responsibility. On the former notion, some agents are collectively responsibility for some action(s) just if each of them is responsible. On the latter, some agents are collectively responsible for some action(s) just if they constitute some complex entity which is responsible. Chapter 3 discusses a third possible notion of collective responsibility; collective responsibility as plural responsibility. On this notion, if some agents are collectively responsible for some action(s) it does not follow that any one thing is responsible. In Chapter 4 I present some cases in which collective responsibility as plural responsibility might occur.

² See Section 1.1.1 of this thesis.

Responsibility and Plurality

In this chapter I outline the central problem with which this thesis will be concerned, and some preliminaries which will be relied upon in later chapters. In Section 1.1 I isolate the form of responsibility on which I will focus, before sketching an account of this type of responsibility in Section 1.2. Section 1.3 introduces what I call the problem of collective responsibility. Responding to this problem will be the main concern of later chapters. Finally in Section 1.4 I outline some recent work in the logic and metaphysics of plurals.

1.1. Responsibility

Responsibility is a polysemous concept. A parent might bear responsibility for her child. An official may accept full responsibility for some governmental misdemeanour. A programme might name the drama teacher as the person responsible for the wonderful (or dreadful) school play. Each of these cases seems to exemplify a different sense in which the term ‘responsibility’ can be employed. In this thesis, I will be exclusively concerned with what Thomas Scanlon (1998, p. 248) has called *attributable* responsibility. An agent is *attributably* responsible for an action just if that action is a basis for moral appraisal of her; that is, she may be the object of moral praise, criticism or indifference on the basis of that action. We might otherwise call this form of responsibility ‘moral responsibility’.

I take this type of responsibility to be very familiar. However it often occurs alongside other moral phenomena. In the following two sections I distinguish *attributable* responsibility from two other closely related notions: substantive responsibility and the practice of holding responsible. Whilst these phenomena often coincide with instances of *attributable* responsibility, they are, I suggest, distinct.

1.1.1. Substantive responsibility

Scanlon (1998) distinguishes attributable responsibility from what he calls *substantive* responsibility. For Scanlon, an agent is substantively responsible for an action just if it is appropriate to make her bear the cost of that action. As such, judgements of substantive responsibility are judgements about what agents owe to one another (Scanlon 1998, p. 290). An agent being attributable responsible for some action is neither necessary nor sufficient for her being substantively responsible for it. An agent can be open to moral appraisal on the basis of some action without it being appropriate to make her bear the full costs of it. For instance, a drug addict may be open to moral criticism for having developed an addiction in the first place. However it does not follow from this that it is appropriate to make him alone bear the full costs of his addiction (Scanlon 1998, p. 292). Others may owe it to him to help him bear these costs. Moreover it may be appropriate to make an agent bear the costs of some action even if it is not a basis for moral appraisal of her. For instance, the state might owe it to the drug addict to provide him with rehabilitation treatment, even if it cannot be morally criticised for his addiction.

To further illustrate the distinction between attributable and substantive responsibility, consider the example of tort law.³ Tort law is concerned with cases in which one party (the defendant) is claimed to have wrongfully gained at the expense of, or at least wrongfully injured, another (the plaintiff). If it is established that the defendant did in fact wrong the plaintiff, we may say that the defendant is responsible for the plaintiff's injury. In an attributable sense, this judgement expresses that the plaintiff's injury is the result of an action which is a basis for moral criticism of the defendant. This is certainly a part of what judgements of responsibility in tort law are supposed to establish: whether or not the defendant has acted wrongfully. In a substantive sense, the judgement of responsibility says that the defendant owes it to the plaintiff to bear the costs of her injury. This expresses a further important aspect of tort law: the payment of damages. When the defendant has wrongfully injured the plaintiff, she is liable to pay damages to the plaintiff as compensation for the

³ For helpful philosophical discussions of tort law see e.g. (Stone 2001; Weinrib 2013).

loss. These two aspects of tort responsibility may come apart. For instance, a defendant may be unable to pay full damages and as such it may be judged inappropriate for her to bear the full cost of the plaintiff's injury. Nevertheless, it does not follow that the defendant is not open to moral criticism on the basis of her action.

1.1.2. Holding responsible

We may also distinguish between an agent being attributably responsible for an action, and her being *held* responsible for that action. The former consists in an agent being open to certain moral judgements on the basis of an action. The latter consists in someone (either another or herself) forming such judgements about the agent and responding to her in particular ways. That is, an agent being attributably responsible for some action consists in the fact that she has acted contrary to or in accordance with moral principles. By contrast, her being held responsible for that action involves people reacting to her behaviour by treating and thinking about her in certain ways. Such reactions form part of a wider social practice of holding responsible.

Peter Strawson (1962) and, following him, R. Jay Wallace (1994) have provided accounts of the practice of holding people responsible. These accounts focus on certain moral emotions or 'reactive attitudes', such as resentment and indignation. According to Strawson and Wallace, when one believes that someone has acted wrongfully towards oneself or someone else, one will be resentful of them or feel indignant towards them. Feeling these ways about the actions and attitudes of others constitute ways in which we hold others responsible for their actions. That is, experiencing and expressing such emotions are ways in which we think about and treat others when we believe them to have acted wrongfully.

Blaming might also be thought to be part of the practice of holding responsible. When one blames another, one treats them and thinks about them in a certain way in response their having acted in a way which one believes to be morally wrong. The practice of holding responsible may also cover the enforcement of certain legal sanctions. When someone is judged to have acted

contrary to certain laws it may be appropriate to enforce a fine on them, for instance, or restrict some of their civil liberties, or even to incarcerate them. Such legal practices are ways in which people are treated in response to the judgement that they have acted illegally.

The practice of holding people responsible for their actions is importantly connected to the concept of someone being morally, or attributably, responsible. Nevertheless an agent being held responsible for an action is neither necessary nor sufficient for her being attributably responsible for it. For instance, Scanlon (2003, p. 362) imagines a friend who has acted wrongfully and so is open to moral criticism on the basis of that action. However the friend has suffered a great deal recently and you sympathise strongly with her. Scanlon suggests that, depending on the severity of the wrong she has committed, one may not be capable of blaming or feeling indignant towards the friend. Consequently, whilst the friend is responsible for her action, she will not be held responsible for it. Furthermore it is possible, and perhaps justifiable, that someone might be held responsible for an action without being responsible for it. For example, suppose a schoolchild steals the teacher's board pens while she is not looking. Unsure of who committed the felony, the teacher gives the whole class a detention. In this case, only one child is attributably responsible for stealing the pens. However each child is held responsible for the action.⁴

An agent being held responsible for an action is also distinct from her being substantively responsible for it. As discussed above, judgements of substantive responsibility are judgements of what people owe to one another. Holding someone responsible requires no such judgements. For instance, if I blame or resent someone for doing something, I do not necessarily make a judgement about whether they owe it to anyone, or whether anyone owes it to them, to bear the costs of that action. Blame, resentment and indignation are reactions

⁴ The teacher's action is not only intelligible, but may also be justifiable. For instance, a procedure of punishing each class member for one child's misbehavior may help to ensure that the children prevent each other from misbehaving. See (Feinberg 1968, pp. 679-681) for a discussion of similar examples.

we have to the perceived wrongdoing of others. Such reactions are silent on who should bear the costs of wrongdoing.

A legal example might help to further illustrate this distinction. Earlier I used the example of tort law to illustrate the notion of substantive responsibility. In tort law, financial sanctions imposed on a defendant are payed as damages to the plaintiff. I suggested that the judgement that a defendant is liable to pay damages to the plaintiff may be conceived of as a judgment that the defendant is substantively responsible for her action; she owes it to the plaintiff to bear the costs of her action. This judgement may coincide with the judgement that the defendant also acted wrongfully or illegally. By contrast, in criminal law sanctions imposed on a defendant are not used to enrich the plaintiff. As such, unlike in tort, the imposition of a criminal sanction may not be conceived of as a judgement about what a defendant owes to a plaintiff. Rather, a criminal sanction is simply an institutional response to an action which is believed to be illegal. As such, criminal sanctions may be conceived of as ways of holding people responsible for illegal actions. However, unlike tort sanctions, they cannot be conceived of as judgements that an agent is substantively responsible for an action.

1.2. Responsibility as self-governance

I have isolated the type of responsibility with which I will be concerned in this thesis: responsibility as a basis for moral appraisal. In this section I will sketch an account of this form of responsibility which I call *responsibility as self-governance*. According to this account, an agent is responsible for an action just if it reflects her sensitivity to and her capacity to act in accordance with practical reasons. I extract this account primarily from the work of Scanlon (1998; 2003). However it also draws on work on responsibility from Strawson (1962) and Wallace (1994) and work on free action from authors such as Harry Frankfurt (1971) and Gary Watson (1975). Such authors have been primarily concerned with developing an account of moral responsibility which

is compatible with the truth of determinism. In this thesis I remain neutral on this topic.⁵

1.2.1. Reasons

In this section I will briefly sketch a conception of practical reasons. I take the notion of a reason to be primitive. A reason for an agent to do something is simply a fact which counts in favour of her doing it. Conversely, a reason for an agent not to do something is a fact that counts against her doing so (Scanlon 1998, p. 17; Thomson 2001, p. 32). On this view, when I have a reason to perform some action ϕ there is a fact of either the form ‘if I ϕ then ...’ or ‘if I don’t ϕ then ...’ which has some normative force in support of my ϕ -ing. When I have a reason not to ϕ there is a fact of the form ‘if I ϕ then ...’ or ‘if I don’t ϕ then ...’ which has some normative force in opposition to my ϕ -ing. What sorts of fact support or oppose an agent doing something? J. J. Thomson (2001, pp. 37-38) singles out evaluative facts as facts of this sort; facts like ‘my doing that thing would be good for me’, or ‘my doing that thing would violate someone’s rights’, or ‘my doing that thing is the best possible thing I could do’, and so on. I take it as basic that there are evaluative facts which support or oppose the performance of certain actions.

On the view I am adopting, there are two kinds of reason for action: *intelligible* and *conclusive*.⁶ The difference between these types of reason consists in the normative force with which they support or oppose an agent doing something. An intelligible reason merely *recommends* to an agent that she perform, or not perform, some action. For instance, suppose I have been invited to a party and I know that a dish which I am fond of will be served there. In this case, I have an intelligible reason to go to the party. If I go then I will have the pleasure of eating food which I like. This fact recommends to me that I go to the party. By contrast, a conclusive reason *demand*s of an agent that she do, or not do, something. To illustrate, suppose I made a promise to

⁵ For alternative accounts of responsibility see e.g. (Schlick 1939, chap. 7; Smart 1961; Strawson 1994).

⁶ My articulation of this distinction is influenced by the work of Elizabeth Anscombe (1967), Véronique Munoz-Dardé (2005) and Thomas Pink (2004; 2007).

my friend that I would go to the party. In this case, I have a conclusive reason to attend the party. If I don't go then I will break a promise which I have made. This fact demands or requires, rather than merely recommending, of me that I attend the party.

The different forces with which intelligible and conclusive reasons support or oppose an agent doing something are revealed in the differing forms of criticism which become appropriate when an agent fails to act in accordance with some reason (Pink 2004; 2007). When an agent fails to act in accordance with an intelligible reason, it may be appropriate to criticise her for acting less than sensibly or foolishly. For instance, suppose that my only reason for going to the party is that there will be food which I am fond of. Moreover suppose that I have no reason to do anything else. In this case, if I don't attend the party then it might be appropriate to make criticisms such as the following: 'it was silly of you not to go to the party', 'the sensible thing to do would have been for you to go to the party', 'you were a fool to stay at home when you could have enjoyed the food at the party', and so on.

By contrast, when an agent fails to act in accordance with a conclusive reason, then her action (or failure to act) may be a basis for moral criticism. For example, suppose that I have promised a friend that I will go to the party, and as such have a conclusive reason to do so. Suppose further that I have no conclusive reason to do anything else. If I then fail to attend to the party then I will be open not merely to the criticism that I acted unwisely or foolishly. My failure to attend the party may be a basis for *moral* criticism of me. As such, the stronger force with which conclusive reasons support or oppose an agent doing something is reflected in the stronger, more serious criticism licenced by an agent's failure to act in accordance with such reasons.

A distinction is sometimes made between there being a reason for an agent to do something, and an agent having a reason to do something. Thomson (2001, pp. 23-24), for instance, holds that the latter entails the former but not *vice versa*. For Thomson, there is a reason for an agent to do something just if there is a fact which counts in favour of her doing it. But this not sufficient for an

agent having a reason to do something. In addition, she must believe that there is a reason for her to do so. Thomson suggests this as an analysis of the ordinary usage of the English phrases. However she does not take this distinction to entail any substantive claims for moral theory. I will use the phrases ‘there is a reason for an agent to do something’ and ‘an agent has a reason to do something’ interchangeably. Both of these phrases are true just in case there is a fact which counts in favour of an agent doing something.

1.2.2. Self-governance

An agent having a conclusive reason to do, or not do, something is a necessary condition for that action being a basis for moral appraisal of that agent; that is, it is a necessary condition for that agent being attributably responsible for that action. But it is not a sufficient condition. In addition, the action must reflect what Scanlon (2003, p. 368) calls the agent’s “capacity for critically reflective, rational self-governance – ‘critically reflective’ because it involves the ability to reflect and pass judgement upon one’s actions and the thought processes leading up to them; ‘rational’ in the broad sense of involving sensitivity to reasons and the ability to weigh them; ‘self-governance’ because it is a process which makes a difference to how one acts”. For short, I will call this the *capacity for self-governance*. When an agent employs her capacity for self-governance, she reflects on what is recommended or demanded of her in a given situation by the reasons which she is sensitive to. Her capacity involves the ability to criticise or affirm an action that she has performed, on the basis of the reasons she had to perform, or not perform, it. Importantly, her capacity for self-governance also involves an ability to do, or not do, what may be recommended or demanded of her by the reasons to which she is responsive.

Moral appraisal is appropriate when, and only when, an agent exercises her capacity for self-governance in particular ways. As Scanlon (2003, p. 369) puts it: “morality is addressed to people who are assumed to have this general capacity [for self-governance], and it tells them how the capacity should be exercised”. How should this capacity be exercised? Assuming Section 1.2.1’s conception of reasons, we can say that an agent ought to employ her capacity for self-governance in accordance with conclusive reasons. Given this, an

action is a basis for moral appraisal of an agent just if it is evidence of the fact that she has exercised, or failed to exercise, her capacity for self-governance according to the demands of conclusive reasons.

Notice that an action may indicate that an agent has employed her capacity for self-governance whilst failing to be a basis for moral appraisal of her. This is the case if the action reflects that the agent has exercised, or failed to exercise, her self-governance merely in accordance with intelligible reasons. I allow that in such cases there may be a sense in which the agent is responsible for the action. For instance, I might be responsible for the unwise, though morally permissible, investments which lead to my being less well-off than I could have been. However this form of responsibility is not my concern here. My focus is exclusively on responsibility as a basis for moral appraisal.

Henceforth, when I say that an action reflects an agent's self-governance, I mean that it reflects the fact that she has employed, or failed to employ, the capacity in accordance with conclusive reasons.

On the account of responsibility put forward here, an action is a basis for moral criticism of an agent just if it is indicative of a failure of self-governance on her part. Such a failure may occur in broadly two ways (Scanlon 1998, pp. 268-270; 2003, p. 368). An agent may recognise some consideration as a conclusive reason to do, or not do, something but then fail to act in accordance with this reason. Here, the agent is sufficiently sensitive to reasons, but fails to act on the basis of this sensitivity. For instance, suppose a young man with no driving licence or experience acknowledges the fact that if he drives his mother's car on the road he will seriously risk causing harm to others, and recognises this as a conclusive reason not to do so. However he decides to drive the car anyway. In this case the young man has made a conscious decision to ignore a consideration which he himself recognises as a conclusive reason not to drive the car. As such, his action reflects a failure on his part to employ his capacity for self-governance as he should.

Alternatively, an agent may fail to attend to some consideration which is a conclusive reason for her to perform, or not perform, some action. Here, the

agent's failure is not that she has failed to act on the basis of her sensitivity to reasons, but rather that she has failed to be sufficiently sensitive to reasons. For instance, now suppose that the young driver does not recognise that his lack of driving abilities gives him a conclusive reason not to drive the car. There is a party at a friend's house later and the only way he can attend is by taking the car. In his eagerness to attend the party, the driver fails to take seriously enough the high risk of his injuring others by driving. Again, the young man's action reflects a failure on his part to exercise his capacity for self-governance as he should. However, unlike in the previous example, he has not resolved to ignore the demand which this reason makes of him. Rather, he has failed to notice this demand. As such, he has failed to be as sensitive to reasons as he could and should have been.⁷

Responsibility as self-governance not only ascribes responsibility for actions, but can also ascribe responsibility for omissions or failures to act. According to this account, an agent is responsible for the fact that she did not perform some action just if that fact reflects her self-governance. For instance, suppose a road worker, in a hurry to finish the job, forgets to replace the cover over a manhole. Leaving the manhole uncovered poses a serious risk of injury to others. As such, there was a conclusive reason for the road worker to replace the cover. The fact that she did not do so is indicative of the fact that she failed to recognise this consideration, due to her impatience to finish the job. Consequently, the road worker's omission to replace the manhole cover is a basis for moral appraisal of her.

1.3. The problem of collective responsibility

So far I have exclusively discussed a form of responsibility which a single agent bears in relation to an action or omission; the responsibility she bears when an action or omission is a basis for moral appraisal of her. I have sketched what I take to be a plausible account of the conditions under which an agent bears this type of responsibility. However it seems as though it is not just *an agent* which can possess this form of responsibility. *Multiple agents*

⁷ See (Shiffrin n.d.) for further discussion of this form of moral failure.

appear capable of collectively bearing attributable responsibility. That is, it seems possible that an action, or some actions, might be a basis for moral appraisal of several agents. Given this, a question arises; can an account of the conditions under which an agent is responsible for an action provide a general account of the conditions under which several agents are collectively responsible for some action(s)? More specifically, is the notion of collective responsibility consistent with responsibility as self-governance?

I will refer to this question as *the problem of collective responsibility*.

Responding to this problem will be the central concern of this thesis. There are broadly two types of response to this problem. One might have what I will call a *radical* response. On a radical response, several agents being collectively responsible for some action(s) is inconsistent with responsibility as self-governance. If some agents are collectively responsible for some action(s) then responsibility as self-governance will be unable to account for their responsibility. Proponents of a radical response face three possibilities. Either responsibility as self-governance is correct as an account of singular responsibility but not as an account of collective responsibility (in which case an additional theory is required for collective responsibility), responsibility as self-governance is incorrect both as an account of singular and of collective responsibility (in which case the same theory may, or may not, account for both singular and collective responsibility); or collective responsibility is impossible.

Alternatively one might have a *deflationary* response to the problem. On this view, collective responsibility and responsibility as self-governance are consistent. If several agents are collectively responsible for some action(s), then responsibility as self-governance can account for their responsibility. If such a response is true then there is no substantive difference between singular and collective responsibility; the same theory accounts for both. In the remaining chapters of this thesis I argue that a deflationary response is possible.

What elements of responsibility as self-governance might a notion of collective responsibility depart from? Writers on collective responsibility have tended to focus on two claims;

(1) Something is responsible for an action only if it is an agent.

(2) An agent is responsible for an action only if she performed it.

Both claims follow from responsibility as self-governance. I will employ a fairly weak conception of an agent; something is an agent just in case it has a capacity for self-governance. That is, agents are things which can reflect on what is recommended or demanded of them by certain features of a situation, and can act on the basis of this sensitivity to reasons. Given this, it's clear that a non-agent cannot be responsible for an action according to responsibility as self-governance. No action can reflect the self-governance of something which lacks such a capacity. As such, this account of responsibility entails (1).

I suggest that it also entails (2). If an action was performed by someone other than myself, then that action will not reflect my ability to regulate my own behaviour in accordance with reasons. Only actions which I have performed are indicative of my capacity for self-governance. One might question this claim. Sometimes an action can appear to reflect an agent's ability to regulate her own behaviour on the basis of reasons, even when she didn't perform it. For instance, suppose a child who is under my care picks up a pan of water off the stove and is scalded by the boiling contents. Picking up the pan is not an action which I performed: it is the child's action. However one might think that this action reflects a failure of self-governance on my part; it is indicative of my failure to take seriously enough the duty of care I have towards the child. As such, I am responsible for the child's action.

This claim, I suggest, is too hasty. The child's action is not indicative of my self-governance *per se*. Rather, it is indicative of my own omissions to act as I should have done; my failure to stop the child from picking up the pan, or my failure to put the pan somewhere where she couldn't have reached it, and so

on. The fact that I did not do these things is indicative of a failure of self-governance on my part. That is, the child's picking up the pan and scalding herself is evidence of the fact that I didn't do some things which I should have done. It is the latter fact which reflects my self-governance. As such, I am responsible for my omission to act rather than for the child's action.

Certain writers have worried that if several agents are collectively responsible for some action(s), then either (1) or (2) will be false; either some non-agent will be responsible for some action(s), or an agent will be responsible for something she didn't do. If this is the case then a radical response to the problem of collective responsibility will be true. That is, if some agents are collectively responsible for some action(s), then the action(s) will not reflect their self-governance. Alternatively one might think that collective responsibility is consistent with (1) and (2). If this is the case, then some necessary conditions for the truth of a deflationary response will hold.

Some authors have complained that the notion of collective responsibility violates principles of fairness. Stephen Sverdlik (1987, p. 68), for instance, worries that collective responsibility might violate claim (2). This, he argues, would be unfair. Worries of this sort are not my concern here. Principles of fairness are often supposed to govern the practice of holding people responsible.⁸ Such principles do not, however, determine whether someone is or is not responsible for something. For instance, on responsibility as self-governance, the latter is determined by whether an action or omission reflects an agent's self-governance: not by considerations of fairness. Given this, the worry that collective responsibility might be in tension with principles of fairness is a worry about *holding* several agents collectively responsible. It is not a worry about several agents *being* collectively responsible, and whether this is consistent with an account of singular responsibility. The latter is exclusively my concern here.

⁸ See e.g. (Pink 2007; Wallace 1994).

1.4. Plurality

Giving a response to the problem of collective responsibility requires consideration of what collective responsibility is. This complicates the task of responding to the problem. There are, I suggest, various senses in which several agents might be collectively responsible for some action(s). This is so even when we specify that ‘responsible’ is being used to mean attributable responsibility. The reason for this is that there is more than one way in which a plurality of things might possess a property, and more than one way in which a plurality of agents might act. This section introduces some ways of talking about pluralities. I begin by introducing the notion of a plural logic. I then move on to discuss applications of central insights taken from work in plural logic. These discussions will be relied upon in later chapters in which I analyse different conceptions of collective responsibility.

1.4.1. Plural logics

In English, it is common to find phrases which appear to refer to more than one thing at a time. For instance, ‘the books’ seems to denote not one thing, but several things: some books. However Frege-Russell predicate calculus treats this feature of natural language in the same way as it treats, for instance, gender: as a mere grammatical feature of natural language, and not as a feature of its logical form. In classical predicate calculus all reference, quantification and predication is singular. A proper name denotes exactly one object, a variable denotes one object at a time, a predicate is satisfied by one subject at a time, and so on.

This singularity of classical predicate calculus has led to problems in notating certain sentences about pluralities.⁹ For instance, consider the sentence ‘Russell and Whitehead wrote *Principia Mathematica*’. This sentence is true. However classical predicate calculus would typically paraphrase it as ‘Russell wrote *Principia Mathematica* and Whitehead wrote *Principia Mathematica*’, and then notate this paraphrase accordingly. This is because, on classical predicate calculus, the predicate ‘wrote *Principia Mathematica*’ may only be

⁹ See e.g. (Massey 1976; Oliver and Smiley 2001).

satisfied by one thing at a time. But the paraphrase is false; neither Russell nor Whitehead wrote *Principia Mathematica*.

In recent years, various logicians and philosophers have challenged classical predicate calculus' singularity and have developed logics which involve plural reference, predication and quantification.¹⁰ Call such logics *plural logics*. In addition to the terms of predicate calculus, such a logic contains names and variables which refer plurally, and quantifiers which quantify over pluralities of things. It also contains a logical dyadic relation, read as 'is one of', which takes as its left-hand argument a singular name or variable, and as its right-hand argument a plural name or variable. A plural logic may additionally contain 'non-distributive' predicates. A non-distributive predicate holds of a plurality of things only if it does not hold of each one of them. For instance, where 'xx' is a plural variable, 'for some xx' is a plural existential quantifier and 'F' is a non-distributive predicate, 'for some xx, Fxx' is true if and only if there are some xx such that xx satisfy F and, for every x, if x is one of xx then x does not satisfy F. In this thesis I will make use of a convention introduced by George Boolos (1984). This convention stipulates that 'for some xx, Fxx' ('some things are F') may be true if there is just one thing that is F.¹¹

A plural logic containing non-distributive predicates faces no problems in notating sentences like 'Russell and Whitehead wrote *Principia Mathematica*'. The predicate 'wrote *Principia Mathematica*' can be treated as a non-distributive predicate. We can then specify that this predicate is satisfied by the plurality Russell and Whitehead. But, since the predicate is non-distributive, it will be satisfied by no one of this plurality: giving the correct result that Russell and Whitehead wrote *Principia Mathematica* whilst neither Russell nor Whitehead did so.

¹⁰ See e.g. (Boolos 1984; 1985; Hossack 2000; Higginbotham 1998; Linnebo 2003; McKay 2006; Oliver and Smiley 2013; Rayo 2002; Uzquiano 2011).

¹¹ My construal of plural logic follows that of Augustin Rayo (2002) and Øystein Linnebo (2003).

Developments in plural logic provide insights which can be applied in other areas of philosophy. Some of the most prominent applications have occurred in the philosophy of logic and mathematics.¹² The following sections discuss two other applications in the topics of property instantiation and action.

1.4.2. Property possession

According to plural logics, relations such as reference need not be one-to-one. That is, reference is not necessarily a relation between a term and an object. Sometimes the referring relation may hold between a term and several objects, as when ‘the books’ refers to several books. The general insight that certain relations are not necessarily one-to-one, I suggest, has various applications in philosophy. One such application is in the case property instantiation. One might think that property instantiation is a relation which holds between a property, or universal, and an object. However, given insights from plural logic, we can deny that this is necessarily the case. Perhaps sometimes the relation of property instantiation can hold between a property and several objects. Following Thomas Smith (2009; 2011), we can distinguish two ways in which a plurality of objects might possess a property (where ‘ xx ’ is a plural variable and ‘for every xx ’ is a plural universal quantifier);

(S) For every xx , xx are severally F if and only if for every x that is one of xx , there is an instance i of F and i is x ’s F -ness.

(P) For every xx , xx are plurally F if and only if there is an instance i of F and i is xx ’s F -ness.

According to (S), when a plurality possess a property severally, for every one of the plurality there is an instance of the property such that there is a one-to-one relation of property possession which holds between it and the instance of the property. By contrast, on (P), when a plurality possess a property plurally, there is a one-to-many relation between an instance of the property and the plurality. To illustrate, some books might be severally heavy. On (S), this is

¹² See e.g. (Boolos 1984; 1985; Higginbotham 1998; Linnebo 2003; Uzquiano 2011).

the case just when for every one of the books there is an instance of heaviness which it possesses: if, for instance, the books are ten copies of *Principia Mathematica*. The books might also be plurally heavy. On (P), this is the case just when there is an instance of heaviness which the books possess: if, for instance, the books are stacked on top of each other in a heavy pile.

Notice that some things possessing a property severally neither entails nor rules out that they also possess it plurally. For example, suppose some books are severally heavy. It doesn't follow from this either that the books are plurally heavy or that they are not. The books may be scattered across a desk, for instance, in which case they will be severally heavy but not plurally heavy. Alternatively they might be stacked on top of each other, in which case they will be both severally and plurally heavy. Moreover suppose that the books are plurally heavy. This entails neither that they are severally heavy, nor that they are not. The books might be ten copies of *Principia Mathematica*, in which case they will be both plurally and severally heavy. Alternatively they might be fifty copies of Wittgenstein's *Tractatus*, in which case they will be plurally but not severally heavy.

It will be helpful to clarify the relationship between several/plural property possession and the notions of distribution/non-distribution mentioned above in relation to predicates. Some things possess a property F distributively just if each of them is F . By contrast, some things are non-distributively F just if they are F but no one of them is F . Some things possessing a property distributively is necessary and sufficient for them possessing it severally; each of some books is heavy, for example, if and only if for every one the books, there is an instance of heaviness which it possesses. However, as Smith (2009, p. 32) points out, some things possessing a property non-distributively is neither necessary nor sufficient for them possessing it plurally. For instance, a pile of ten copies of *Principia Mathematica* will be plurally heavy: the books will possess an instance of heaviness. However the books will not be non-distributively heavy, since each of them is heavy. So a plurality possessing a property non-distributively is not necessary for them possessing it plurally. Furthermore two piles of fifty copies of the *Tractatus* will be non-

distributively heavy. There is a sense in which the books are heavy, each of the two piles are plurally heavy, and no one of the books is heavy. However the books will not be plurally heavy, since there is no instance of heaviness which is possessed by them. Rather, there is an instance of heaviness possessed by each pile of books. So a plurality possessing a property non-distributively is not sufficient for them possessing a property plurally.

The notion of plural property possession has further application in the topic of ontology. For instance, some writers have argued that plural property possession provides a defence of ontological atomism: the view that there are no complex objects.¹³ Ontological atomists deny that there are such things as tables and chairs, for instance. Such things are not objects which are constituted out of further less complex objects. That is, the relation ‘constitute’, which takes a plural term as its left-hand argument and a singular term as its right-hand argument, is never instantiated. Rather, putative complex objects are identical with pluralities of simple objects, or atoms, arranged in particular ways. However we ordinarily talk about tables and chairs as objects which can possess properties: ‘that table is heavy’, ‘this chair is small’, and so on. Proponents of ontological atomism have argued that, using the notion of plural property possession, such sentences can be paraphrased so that they remain true but involve no reference to complex objects. For instance, ‘the table is heavy’ can be paraphrased as ‘there are some *xx* such that *xx* are arranged table-wise and *xx* are plurally heavy’. In the paraphrase, heaviness is re-construed as a property possessed plurally by some simple objects, rather than a property possessed singularly by some complex object. As such, the paraphrase can be true even if there are no such things as tables.¹⁴

¹³ Fraser MacBride (2005) also has used the notion of plural property possession to attack a Russellian account of the universal-particular distinction.

¹⁴ For discussion of this view see e.g. (Hossack 2000; Uzquiano 2004; van Inwagen 1990, Section 11).

1.4.3. Action

A further application of insights from plural logic is in the case of action. One might think that performing an action can be conceived of as a one-to-one relation between an action (an instance of an act-type) and an agent. However, following insights from plural logic, we can challenge whether this is necessarily the case. Perhaps performing an action can sometimes be conceived of as a one-to-many relation between an action and several agents. Following Smith (2011, p. 216), we can distinguish two ways in which several agents can act (where φ is an act-type);

(SA) For every xx , $xx \varphi$ severally if and only if for every x that is one of xx , there is an instance i of φ and i is x 's φ -ing.

(PA) For every xx , $xx \varphi$ plurally if and only if there is an instance i of φ and i is xx 's φ -ing.

When a plurality of agents act severally, according to (SA), for every agent that is one of that plurality there is an action such that there is a one-to-one relation of action performance which holds between the action and the agent. By contrast, on (PA) when several agents act plurally, there is an action such that there is a one-to-many relation of action performance which holds between the action and the agents. To illustrate, the writing of *The Problems of Philosophy* and the writing of *Process and Reality* are each instances of the act-type 'writing a book'. Moreover each of these actions was performed by a single agent; the former was performed by Russell and the latter by Whitehead. So, on (SA), it's true that Russell and Whitehead wrote a book severally. Similarly, one might think that the writing of *Principia Mathematica* was also an instance of the act-type 'writing a book'. But it seems implausible to think that this action had a single agent; no one wrote *Principia Mathematica* on their own. Given (PA), we can say that the writing of *Principia Mathematica* was an action performed plurally by two agents: Russell and Whitehead.

Given that there is such a thing as acting plurally, there appears to be a relationship between plural actions (actions with several agents) and singular actions (actions with one agent). For instance, when Russell and Whitehead plurally wrote *Principia Mathematica* they also severally performed some singular actions. Russell, let's suppose, wrote some sections, whilst Whitehead wrote others, and so on. There are a number of theories available concerning what a plural action is, and under what conditions such an action might arise.¹⁵ Among other things, such a theory may tell us what the relationship between singular and plural action is. One might think, for instance, that a plural action is a complex event or process which is constituted by some plurality of singular actions. Alternatively one might think that plural actions merely supervene on, or *are*, pluralities of singular actions. In the following chapters I will talk of singular actions as parts of plural actions. By using this terminology I do not intend to commit myself to any metaphysical view on the relationship between singular and plural action.

Different theories of plural action will be more or less restrictive as to what counts as a plural action. Michael Bratman (1992), for instance, suggests that you and I ϕ plurally only if I intend that we ϕ , you intend that we ϕ and this is common knowledge between us. By contrast, Frank Jackson (1987) appears to have a theory of plural action according to which any mereological sum of singular actions is a plural action.¹⁶ Various pluralities of singular actions will count as plural actions on Jackson's account, but not on Bratman's. In this thesis I remain neutral as to which, if any, of the available theories of plural action is most plausible.

1.5. Conclusion

This chapter has introduced various preliminaries on which I will rely in articulating the central argument of this thesis. I have distinguished attributable responsibility from various other forms of responsibility.

¹⁵ See e.g. (Bratman 1992; Gilbert 2013; Gold and Sugden 2007; Kutz 2000, chap. 3; Miller and Tuomela 1988; Searle 1990; Smith 2011; Velleman 1997).

¹⁶ Jackson himself may not propose an account of plural action *per se*. However such a theory can be developed along the lines of his discussion. See note 22 for further discussion.

Moreover I have sketched an account of this form of responsibility, according to which an agent is attributably responsible for an action just if it reflects her self-governance. Given this, we can pose the problem of collective responsibility; is the notion of collective responsibility consistent with responsibility as self-governance? I concluded this chapter with a discussion of plural logics, plural property possession and plural action. The notions introduced in this discussion will be employed in the following chapters, in which I will distinguish three notions of collective responsibility corresponding to three ways in which several things might possess a property.

Several and Corporate Responsibility

In this chapter I discuss two forms of collective responsibility; collective responsibility as several responsibility and collective responsibility as corporate responsibility. On the former notion, some agents are collectively responsible for some action(s) just if each of them is responsible. On the latter, some agents are collectively responsible for some action(s) just if they constitute some complex entity which is responsible. Section 2.1 discusses collective responsibility as several responsibility, whilst Section 2.2 is devoted to collective responsibility as corporate responsibility. These notions of collective responsibility correspond to two ways in which several things might possess a property; by each of them possessing an instance of the property, or by those things constituting some further entity which possesses an instance of the property. This chapter argues that both notions of collective responsibility are compatible with responsibility as self-governance.

2.1. Several responsibility

In this section I will focus on a form of collective responsibility which I will call *collective responsibility as several responsibility*. In this sense, when several agents are collectively responsible for some action(s), the agents are severally open to moral appraisal on the basis of the action(s); that is, given Section 1.4.2's (S), each of the agents is responsible for some action(s).

Given this conception of collective responsibility, we can be more precise about what the possible responses to the problem of collective responsibility will say. Collective responsibility as several responsibility says that if some agents are collectively responsible for some action(s) then each of them is open to moral appraisal on the basis of some action(s). Given this, a deflationary response says that if several agents are collectively responsible for some action(s), then the action(s) reflect(s) each agent's capacity for self-

governance. By contrast, on a radical response, if several agents are collectively responsible for some action(s), then it is not the case that the action(s) reflect(s) each agent's self-governance.

2.1.1. Lewis' challenge

H. D. Lewis (1948) thought that the notion of collective responsibility violated Section 1.3's claim (2). That is, according to Lewis, if several agents are collectively responsible for some action(s), then at least one of them is responsible for something she didn't do. If Lewis is correct then a radical response to the problem of collective response is true. What would collective responsibility have to be in order for this position to be correct? I suggest that Lewis might have had something like the following conception in mind;

(SR1) For every plurality of agents xx , xx are collectively responsible for some action(s) aa if and only if for every x that is one of xx and for every a that one of aa , x is responsible for a .

According to (SR1), when a plurality of agents are collectively responsible for some action(s), each of the agents is responsible for each action. Why think that this notion of collective responsibility involves someone being responsible for something she didn't do? Lewis seems to adopt a further assumption which I will call *individualism about action*; for every action, there is exactly one agent who performed it. This assumption entails that there is no such thing as plural action, as elaborated by Section 1.4.3's (PA). Given (SR1) and individualism about action, we can give the following argument to establish Lewis' conclusion;

Premise 1: For some agents xx , xx are collectively responsible for some action(s) aa .

Premise 2: From (SR1) – for every x that is one of xx and for every a that is one of aa , x is responsible for a .

Premise 3: From individualism about action – for every a that is one of aa , there is exactly one agent who performed a .

Conclusion: From premises 2 and 3 – for every x that is one of xx and for some a that is one of aa , x is responsible for a and x did not perform a .

To illustrate this result, consider the following variant of an example discussed by Lewis (1948, pp. 12-13). Several burglars commit a burglary; one breaks the lock on the door, one rushes in and takes the valuables, one stands guard outside, one drives the getaway car, and so on. Given Lewis' individualism about action, it follows that there is not just one instance of burglary performed by the burglars. That is, the burglars do not burgle plurally. Rather there are several actions; an instance of lock-breaking, some instances of possession-grabbing, an instance of guarding, and so on. Each of these actions has exactly one of the burglars as its agent. It is this plurality of actions which gets referred to as 'the burglary'. In Lewis' words (1948, p. 13): "it would be a great mistake to suppose that we have here a single criminal operation...it is not one crime that we have but many". Given this, if the burglars are collectively responsible for the burglary then, on (SR1), each one of them is responsible for each of the above actions. However this contradicts (2). For instance, it follows that the lock-breaker will be responsible not only for the instance of lock-breaking, an action which she performed, but also for several actions which she did not perform; the instances of possession-grabbing, the instance guarding, and so on.

Lewis takes this to show that collective responsibility is not possible. My reconstructed argument reveals that, given (SR1) and individualism about action, some agents being collectively responsible for some action(s) is inconsistent with (2), and as such is inconsistent with responsibility as self-governance. Consequently if collective responsibility is possible then there must be an alternative form of moral responsibility: one that does not entail claim (2). Lewis (1948, p. 3) is "especially certain" of the principle that "no one can be responsible, in the properly ethical sense, for the conduct of another". However if there is reason to doubt claim (2) then, for Lewis (1948, p. 3), "the proper procedure will be, not to revert to the barbarous notion of collective or group responsibility, but to give up altogether the view that we are accountable in any distinctively moral sense". That is, any reason to doubt

(2) is not reason to think that collective responsibility might be possible, but rather reason to doubt that moral responsibility is possible.

2.1.2. Responses to Lewis

How might one respond to Lewis' challenge? There are two related ways in which one might challenge the reconstructed argument which I have attributed to Lewis. One might think that there are forms of collective responsibility other than (SR1). Alternatively one might deny Lewis' individualism about action and hold that there are plural actions. In this section I explore both options, limiting alternatives to (SR1) to conceptions of collective responsibility as several responsibility.

An alternative to (SR1): Let's first explore the former strategy. How else might one understand collective responsibility as several responsibility, other than as (SR1)? One might suggest the following;

(SR2) For every plurality of agents xx , xx are collectively responsible for some actions aa if and only if for every x that is one of xx , there is an a that is one of aa such that x performed a and x is responsible for a .

Unlike (SR1), according to (SR2) when some agents are collectively responsible for some actions, each one of those agents is responsible for an action that she performed. As such, on (SR2), collective responsibility is consistent with (2), even given Lewis' individualism about action. As such, conceiving of collective responsibility as (SR2) allows for a deflationary response to the problem of collective responsibility. That is, when several agents are collectively responsible for some actions, for each agent there is an action that reflects her self-governance.

To illustrate, consider again the case of the burglars. On (SR2), if the burglars are collectively responsible for the burglary, then each burglar is responsible only for the action(s) which she performed. The lock-breaker, for instance, will be responsible only for the instance of lock-breaking, the guard will be responsible only for the instance of guarding, and so on. Given this, one need

not appeal to a theory other than responsibility as self-governance in order to account for the burglars' responsibility. For instance, since the instance of lock-breaking was an action that the lock-breaker performed, it may reflect her self-governance.

(SR2) can be modified such that it allows that several agents may be collectively responsible for some failure(s) to act. So modified, (SR2) will say that some agents are collectively responsible for some omission(s) to act just if each of those agents is responsible for the fact that she herself did not act. To illustrate, consider the following example adapted from Joel Feinberg (1968, pp. 683-684). A swimmer finds himself in trouble off a public beach. There is no professional life guard on duty, however there are plenty of competent swimmers sunbathing on the beach. Each of the sunbathers recognises that the swimmer is in trouble, and is capable of saving him without putting herself at excessive risk. However nobody comes to his aid. As a result, the swimmer drowns.

The modified (SR2) says that the sunbathers are collectively responsible for their inaction just if each sunbather is responsible for the fact that she herself did not save the swimmer. This is consistent with responsibility as self-governance. The fact that a sunbather did not save the swimmer may reflect a failure on her part to regulate her behaviour in accordance with conclusive reasons. By contrast a modified (SR1), tailored to discussion of inaction rather than action, will say that the sunbathers are collectively responsible for their inaction just if each sunbather is responsible for the fact that every sunbather failed to save the swimmer. This is inconsistent with responsibility as self-governance. The fact that a sunbather did not save the swimmer cannot reflect another sunbather's capacity for self-governance.

So, unlike (SR1), (SR2) allows for a deflationary response to the problem of collective responsibility. Furthermore something like (SR2) is undoubtedly correct as a reading of the sentence 'some agents are responsible for some things'. It seems intuitively correct, for instance, to say of the sunbathers that they are responsible for not saving the drowning swimmer, and a version of

(SR2) gives an appealing reading of this claim; each one of the sunbathers is responsible for her own inaction (Feinberg 1968, pp. 683-684).

However as a reading of the sentence “some agents are *collectively* responsible for some things” I suggest that (SR2) is less plausible. Writers who are sceptical about the possibility, or at least the attractiveness, of collective responsibility seem happy to allow that several agents might be responsible for some things in the sense provided by (SR2). Lewis (1948, p. 12-13), for instance, claims that each of the burglars in his example is responsible for whatever action(s) she performed in the course of the burglary. Nevertheless he denies that the burglars are collectively responsible for the burglary.¹⁷

Furthermore even if one is not sceptical about collective responsibility, (SR2) seems to over-generate cases of it. For instance, on (SR2), Nathuram Godse and James Earl Ray were collectively responsible for the assassinations of Mahatma Ghandi and Martin Luther King Jr; the assassination of Ghandi was an action performed by Godse for which he was responsible, and the assassination of King was an action performed by Ray for which he was responsible. However this result seems implausible. Ray assassinated King in 1968: 18 years after Godse was executed for the assassination of Ghandi. It seems, at best, jarring to think that Ray and Godse could have been collectively responsible for these two assassinations.¹⁸

Responsibility for plural actions: Are there any further conceptions of collective responsibility as several responsibility? I suggest that an alternative to both (SR1) and (SR2) is available if, *contra* Lewis, one accepts that several

¹⁷ See (Narveson 2002) for a similar position.

¹⁸ In her response to Feinberg’s paper, Hannah Arendt (1987, p. 44) seems to have a similar reaction to his swimming case: “No *collective* [Arendt’s emphasis] responsibility is involved in the case of the thousand experienced swimmers, lolling at a public beach and letting a man drown in the sea without coming to his help, because they were no collectivity to begin with”. I am sympathetic to the view that Feinberg’s sunbathers are not collectively responsible for their inaction; though they are responsible for it, in the sense that each is responsible for her own inaction.

agents can act plurally. This assumption allows for the following conception of collective responsibility as several responsibility;

(SR3) For every plurality of agents xx , xx are collectively responsible for some action(s) aa if and only if for every x that is one of xx and for every a that is one of aa , x is responsible for a and a was performed plurally by xx .

According to (SR3), some agents may be collectively responsible for some action(s) only if the action(s) was/were performed plurally by the agents. Like (SR2) and unlike (SR1), I suggest that such a conception allows for a deflationary response to the problem of collective responsibility.

To illustrate, consider again Lewis' burglary example. On (SR3), the burglars are collectively responsible for the burglary just if the burglary is an action performed plurally by the burglars (or is a plurality of actions, each of which was performed plurally by the burglars), and each burglar is responsible for that plural action. Rejecting individualism about action, one may posit, *contra* Lewis, that the burglary was a single action; an action which has the burglars as its agents. According to (SR3), the burglars are collectively responsible for that action just if each of them is responsible for it. As such, on (SR3), if the burglars are collectively responsible for the burglary, it does not follow that claim (2) is false; that is, it does not follow that any of them is responsible for something she didn't do. Each burglar is responsible for an action of which they are an agent. However they are not the sole agent of the action. Rather, the action has several agents and each burglar is one of those agents.

Given this, a necessary condition for (SR3) being consistent with responsibility as self-governance is fulfilled. However this is not sufficient. In addition, it must be the case that a plural action may reflect an agent's capacity for self-governance. I suggest that this is the case. A plural action may reflect one of its agent's self-governance in broadly two ways. An agent may perform a singular action which reflects her self-governance, and that action may be a part of some plural action. As such, we might say that the plural action also reflects her self-governance. For example, when the lock-breaker breaks the

lock, she performs a singular which reflects a failure of self-governance on her part. Furthermore we can suppose that her breaking the lock is a part of the burglary: a plural action. Given this, there is, I suggest, a sense in which the plural action is indicative of the lock-breaker's failure of self-governance. In 'playing her part' in the burglars' plural act of burglary, the lock-breaker performs a singular action which reflects a failure on her part to employ her capacity for self-governance as she should.

Alternatively, an agent might perform a singular action which, *qua* singular action, reflects no failure of self-governance on her part. However that action may be a part of a plural action which does reflect such a failure. To illustrate, suppose a young man signs up for a role in a prison camp. The purpose of the camp is to detain and punish anyone who opposes the dominant political regime. The young man is aware of this and enthusiastically signed up by his own choice. The man's assigned role is to guard the inmates who labour in the camp, ensuring that they behave according to the camp's regulations. In carrying out this role, the guard does not perform any singular actions which reflect any failure of self-governance on his part. He never harshly disciplines anyone, for instance; the inmates are all submissive and obey the camp's regulations. However we may suppose that, in enacting and maintaining the camp's regulations, the prison staff perform some plural actions; they plurally detain and punish opponents of the regime. Whilst the guard's singular actions reflect no failure of self-governance on his part, one might think that these plural actions do. In enacting his assigned role, the guard intends his own actions to further the end of the camp. That is, he knowingly and willingly 'plays his part' in the prison staff's plurally detaining and punishing opponents of the regime. As such, such plural actions reflect the guard's failure to regulate his own behaviour in accordance with conclusive reasons.¹⁹

(SR3) may also be modified such that it says that several agents may be collectively responsible for some failure(s) to perform some plural action(s). Such a variant will say that some agents are collectively responsible for some

¹⁹ See (Kutz 2000) for further discussion of this form of responsibility for plural actions.

omission(s) to act plurally just if each of them is responsible for the fact that the agents did not act plurally. To illustrate, consider the following example adapted from Virginia Held (1970, pp. 476-477). Several passengers are travelling in a tube carriage. Suddenly, one of the passengers pushes another to the ground and starts to beat her. Whilst a single passenger acting by herself cannot prevent this, the passengers are capable of plurally overpowering the attacker and saving victim without excessive risk to themselves. However the passengers do nothing to prevent the attack, and simply look away. By the time the train reaches the next stop and the attacker has fled the scene, the victim has sustained fatal injuries.

Modified (SR3) says that the passengers are collectively responsible for their inaction just if each is responsible for the fact the passengers did not plurally overpower the attacker. This is consistent with responsibility as self-governance just if this fact may reflect the self-governance of each passenger. It seems plausible, I suggest, that this condition is fulfilled in the above example. Each passenger has a reason to try and coordinate with the others so as to overpower the attacker with them; if she does so then she will increase the likelihood of the victim being saved. The fact that the passengers did not plurally overpower the attacker indicates that each of them failed to regulate their behaviour in accordance with such reasons.

As a conception of collective responsibility as several responsibility, I suggest that (SR3) is preferable to both (SR1) and (SR2). Unlike (SR1), it allows for a deflationary response to the problem of collective responsibility. Furthermore, unlike (SR2), (SR3) need not over-generate cases of collective responsibility. Whether it does so or not depends on the theory of plural action one adopts. One might follow Bratman (1992), for instance, and hold that an action is a plural action just if its agents have some set of interlocking intentions. Given this, (SR3) does not entail that, for instance, Godse and Ray were collectively responsible for the assassinations of Ghandi and King. On a Bratman-style view, neither assassination was performed plurally by Godse and Ray. As a result, on (SR3), they cannot be collectively responsible for either of these actions. By contrast, one might follow Jackson (1987), for instance, and hold

that any mereological sum of singular actions is a plural action. On such a view, the mereological sum of Ghandi's assassination and King's assassination is a plural action, performed by Godse and Ray. Given this, on (SR3), they may be collectively responsible for this plural action.

Unlike (SR1) and (SR2), (SR3) is ontologically committing insofar as it assumes the existence of plural actions. That is, if no plurality of agents can perform an action plurally, then no plurality of agents may be collectively responsible for some action(s) in the sense provided by (SR3). However a commitment to the existence of plural actions is not, I suggest, a particularly strong or controversial commitment to have. As discussed above, various authors have defended the position that there are plural actions, and have provided theories of what they are and how they arise. Moreover there are locutions of ordinary English which appear to reveal a common-sense commitment to the existence of plural actions; "we lifted the table together", "the film was the result of a joint effort", "Russell and Whitehead wrote *Principia Mathematica*", and so on.

2.2. Corporate responsibility

I will now move on from collective responsibility as several responsibility to discuss a different form of collective responsibility: *collective responsibility as corporate responsibility*. On this conception, collective responsibility may be possessed by some agents only if they constitute some further complex entity. The agents are collectively responsible for some action(s) just if that entity is responsible for the action(s). That is, strictly speaking, responsibility is not possessed by a plurality of agents, but rather by a complex entity made up of those agents.

This notion of collective responsibility is most common in law where, under certain conditions, several individuals may be regarded as a 'body corporate' (hence my term 'corporate responsibility'); something with the legal standing of a person. Different legal traditions vary on whether a body corporate is considered as an entity in its own right, or whether it is merely treated under law as if it were an entity (French 1979, pp. 208-210). On the notion I will be

discussing in this section, collective responsibility as corporate responsibility regards body corporates as complex entities. On this view, a body corporate is not merely a legal fiction, but exists as a part of the ontological order.²⁰ To clarify, I am not (yet) restrictive as to what kind of complex entity a body corporate might be. For instance, a body corporate might be a set of agents, or any mereological sum agents.

Given this notion of collective responsibility, we can clarify the possible responses to the problem of collective responsibility. On collective responsibility as corporate responsibility, collective responsibility involves some complex entity being responsible for some action(s). Given this, a deflationary response to the problem of collective responsibility will say that the complex entity being responsible for some action(s) is consistent with the account of responsibility as self-governance. On this position, if several agents are collectively responsible for some action(s), then the action(s) may reflect(s) the self-governance of a complex entity which those agents constitute. By contrast, a radical response will say that responsibility as self-governance cannot give the result that the complex entity is responsible for some action(s). On this view, if collective responsibility is possible, then an account of the responsibility possessed by the complex entity must be provided by an alternative theory.

2.2.1. Collective agents

I believe that a deflationary response is possible on a conception of collective responsibility as corporate responsibility. In order for such a response to be correct, this conception of collective responsibility must be consistent with claims (1) and (2); that is, on this conception, if several agents are collectively responsible for some action(s) it cannot be the case either that something without the capacity for self-governance is responsible, or that an agent is responsible for something she didn't do. In what follows I will argue that these conditions can be fulfilled. In doing so I will place the following restriction on the kind of complex entity a body corporate might be. In order to be

²⁰ For discussion of this view see e.g. (Descombes 2011, chap. 5; French 1982; Smith 2005)

collectively responsible for some action(s), some agents must constitute a *collective agent*: a complex entity which meets the conditions for agency. Employing this notion, we can provide the following precisification of collective responsibility as corporate responsibility;

(CR) For every plurality of agents xx , xx are collectively responsible for some action(s) aa if and only if xx constitute a collective agent c and c is responsible for aa .

Various authors have argued that several agents can constitute a collective agent, and have provided theories of the conditions under which this might occur.²¹ According to such authors, collective agents can perform actions. Given this, (CR) can be consistent with (2); we simply specify that a collective agent may be responsible only for actions which it has performed. It follows that if a collective agent is responsible for an action, it is not responsible for something it did not do.

In order for (CR) to be consistent with (1), a collective agent must be something with a capacity for self-governance. Can this condition be fulfilled? I suggest that it can. Several theorists appear to hold that in order for several agents to constitute a collective agent, there must be some system of rules in place which regulates their behaviour in accordance with certain considerations. Such a system, I suggest, may be considered a functional approximate of a capacity for self-governance.

To illustrate, consider Peter French's (1979) influential account of collective agency. French argues that corporations, such as BP, Sainsbury's and British Gas, are collective agents which are capable of performing actions. According to his account, several agents constitute a collective agent only if they are regulated by what French (1979, pp. 211-213) calls a Corporation's Internal Decision Structure (CID Structure). A CID Structure is a set of rules which

²¹ See e.g. (French 1979; Lawford-Smith 2015; List and Pettit 2011; Pettit and Schweikard 2006)

has two main functions. Firstly, it provides the hierarchy within a corporation's power structure. Secondly, it governs decisions made on behalf of the corporation. This second function of a CID Structure, I suggest, bears important similarities to a capacity for self-governance. As discussed in Section 1.2.2, a capacity for self-governance enables an agent to reflect on what is recommended or required of her by certain reasons, and to act, or not act, in accordance with them. For French (1979, p. 213), a CID Structure contains certain policies which are designed to ensure that certain decisions and actions within the corporation are made or performed for reasons which are consistent with the central policies of the corporation, such as increasing profit. As such, given French's claim that a corporation can perform actions, a CID Structure regulates the actions a corporation performs on the basis of reasons. In this way, a CID Structure appears to be functionally approximate to a corporation's capacity for self-governance.

It therefore appears that (CR) can meet two necessary conditions for being consistent with responsibility as self-governance. A collective agent may be something with a capacity for self-governance. Moreover if a collective agent is responsible for some action(s), it need not be the case that anyone is responsible for something she didn't do. Can an action performed by a collective agent reflect its capacity for self-governance? I see no reason why not. I have suggested that a collective agent's capacity for self-governance may be a set of policies governing certain decisions made by its members. As such, a collective agent's action will reflect its self-governance just if it is indicative of those policies, and the reasons which they highlight for consideration. To illustrate, consider an example from French (1979, pp. 213-215) in which Gulf Oil joins a uranium cartel. According to French, joining the cartel is an action which can be ascribed to Gulf Oil. Given this, I suggest that this action may reflect Gulf Oil's capacity for self-governance. It may, for instance, reflect the fact that the policies regulating decision making within the corporation fail to adequately take into account reasons which demand of the corporation that it not join the cartel.

Insofar as (CR) is consistent with responsibility as self-governance, I suggest that the notion of collective responsibility as corporate responsibility is plausible. However the assumption that there are collective agents which can perform actions and which possess a capacity for self-governance is strong and controversial. To clarify, I am not committed to this assumption. Rather, I have argued that there are accounts of collective agency which appear to be committed to this assumption. Given such accounts, (CR) is consistent with responsibility as self-governance. I accept that these accounts might turn out to be implausible. If so then collective responsibility as corporate responsibility will not be consistent with responsibility as self-governance.

2.2.2. Corporate and several responsibility

It may be helpful to clarify the relationship between collective responsibility as corporate responsibility and collective responsibility as several responsibility. Collective agents are made up of simpler agents, and a collective agent acts by virtue of its members performing actions. As such, it is reasonable to think that there will be some relationship between the responsibility of a collective agent and the responsibility of its members. In light of this, one might think that (CR) is reducible to some form of collective responsibility as several responsibility. That is, one might think that a collective agent is responsible for some action(s) just if each of its members is responsible for the action(s).

I suggest that reductive strategy is implausible. I have demonstrated in Section 2.1.1 that (SR1) may be incompatible with responsibility as self-governance, and as such is less plausible than other conceptions of collective responsibility. So if (CR) is reduced to some form of collective responsibility as several responsibility and its consistency with responsibility as self-governance retained, then it must be reduced to either (SR2) or (SR3). I suggest that neither of these strategies succeeds. Suppose that some collective agent *c* is responsible for some action *a*. If (CR) is reducible to (SR2) then *a* must be identical with some plurality of actions performed severally by the members of *c*. Alternatively, if (CR) is reducible to (SR3) then *a* must be identical with some action(s) performed plurally by *c*'s members. In Section 2.2.1 I

stipulated that a collective agent is responsible for an action only if it performed the action, so as to ensure that (CR) does not violate claim (2). Given this, it follows that *c* performed *a*. But if *a* was performed by *c*, then it cannot have been performed by the members of *c*: either severally or plurally.

To illustrate, consider French's example of Gulf Oil joining a uranium cartel. Suppose there is a company-wide vote on whether or not Gulf Oil joins the cartel, which the 'yes' vote wins. In this case, a plurality of actions performed severally by the corporation's members has determined that the corporation will join the cartel. However it does not follow that joining the cartel is identical with this plurality of actions (French 1979, p. 214). If joining the cartel is an action performed by the collective agent Gulf Oil, then it will not be identical with the plurality of casted votes. This is consistent with the votes being necessary for it being the case the Gulf Oil joins the cartel.

Alternatively one might think that the vote was an action performed plurally by the corporation's members. However it does not follow that this plural action is identical with Gulf Oil's action of joining the cartel. This is because, given Section 1.4.3's (PA), plural actions are not identical to the actions of collective agents. On (PA), a plural action is an action which has several agents. By contrast, an action performed by a collective agent has one agent. So whilst it might be the case the Gulf Oil joined the cartel by virtue of its members acting plurally, it is not that case that those members performed the action of joining the cartel. Given that joining the cartel was performed by a collective agent, it cannot also be the case that this was a plural action performed by several agents.²²

The above demonstrates that if (CR) is reducible to collective responsibility as several responsibility, then (CR) will be inconsistent with (2). If *a* was

²² Various authors do not distinguish between plural actions and the actions of collective agents e.g. (Collins 2011 p. 231; Corlett 2001, p. 575). Smith (2009, p. 36) also identifies this lack of distinction in Jackson (1987). Equating plural actions with the actions of collective agents is an instance of a more general strategy to avoid positing plural states of affairs. This strategy treats the subjects of a plural state of affairs as a single compound object. See (MacBride 2005, pp. 580-584; Oliver and Smiley 2001) for further discussion.

performed by *c*, then *a* was neither performed by a member of *c* nor performed plurally by some members of *c*. Given this, it follows that neither (SR2) nor (SR3) will give the result that the members of *c* are collectively responsible for *a*. (CR) is reducible to either (SR2) or (SR3) only if a collective agent is only ever responsible for an action it didn't perform: some action(s) performed either severally or plurally by its members. But then (CR) violates (2). The only other alternative is to suppose that (CR) is reducible to (SR1), in which case (CR) again violates (2).

To clarify, I do not deny that it may be possible to paraphrase attributions of responsibility to a complex entity as attributions of responsibility to its members. One might resist the assumption that there are collective agents which can perform actions and bear responsibility. If so then one will deny that (CR) provides a plausible conception of collective responsibility. My aim in this section has not been to argue that it is impossible to talk about the responsibility of a body corporate as responsibility born by exclusively by its members.²³ Rather I have argued that, given that (CR) is a plausible conception of collective responsibility, (CR) is not reducible to collective responsibility as several responsibility.

2.2.3. Corporate responsibility and non-agents

I have argued that, on (CR), if several agents are collectively responsible for some action(s), it does not follow that any one of them is responsible for the action(s). Recently, some authors have argued that (CR) is the only plausible conception of collective responsibility which has this consequence. According to such authors, if some agents who do not constitute a collective agent are collectively responsible for some action(s) and no one of them is responsible for the action(s), then claim (1) must be false; that is, it follows that something which is not an agent is responsible for some action(s).²⁴ If so then the conception of collective responsibility employed will be inconsistent with

²³ See (Thompson 2004) for a discussion which treats responsibility within organised, bureaucratic institutions as responsibility born by their members.

²⁴ Stephanie Collins (2013, p. 231) and Holly Lawford-Smith (2015) appear to propose versions of this argument. It is also discussed by Felix Pinkert (2014, pp. 188-190) and Bill Wringe (2012).

responsibility as self-governance, and as such will provide a radical response to the problem of collective responsibility.

Why think that this is the case? The argument relied upon by proponents of this position appears to be along the following lines;

Premise 1: For every x , x is responsible for an action only if x is an agent.

Premise 2: For every xx , xx are collectively responsible for an action a if and only if either (for every x that is one of xx , x is responsible for a) or (there is a y such that xx constitute y and y is responsible for a).

Premise 3: For some xx , xx are collectively responsible for an action a and (there is no x such that x is one of xx and x is responsible for a) and (there is no y such that xx constitute y and y is an agent).

Premise 4: From Premises 1 and 3 – for some xx , xx are collectively responsible for an action a and (there is no x such that x is one of xx and x is responsible for a) and (there is no y such that xx constitute y and y is responsible for a).

Conclusion: From Premises 2 and 4 – for some xx , xx are collectively responsible for an action a and xx are not collectively responsible for a .

The conclusion's first conjunct, that some things are collectively responsible for an action, follows from Premise 4. The second conjunct, that they are also not collectively responsible for the same action, follows from Premises 2 and 4; Premise 2 gives a necessary and sufficient condition for some things being collectively responsible for an action and Premise 4 stipulates that the things in question fail to meet this condition. Since the conclusion is false and the argument is valid, at least one of the premises must be false. Typically, those who rely on this argument wish to reject Premise 4: the claim that some things may be responsible for an action whilst it is neither the case that each of them is responsible for the action, nor that they constitute an entity which is responsible for the action. If Premise 4 is false then at least one of Premises 1 and 3 must be false, since it is entailed by these two claims. Premise 1 is very plausible and is entailed by the account of responsibility as self-governance. Consequently, the argument is taken to be a *reductio* of Premise 3.

However it is not clear that this is the strategy to adopt. Those writers who discuss versions of this argument tend to implicitly assume some variant on Premise 2.²⁵ But it is not obvious that this premise is true, and if it is not then Premises 1, 3 and 4 may be correct. To see this, consider again Section 1.4.2's brief discussion of ontological atomism. In this section I discussed a particular objection to atomism. A version of this objection notes that putative complex objects seem to possess properties which the atoms that putatively constitute such objects do not. For instance, a table might be heavy whilst it is not the case that each atom which putatively constitutes the table is heavy. It then assumes that the table is heavy just if either the table is identical with a plurality of atoms and each of them is heavy, or the table is a complex object and it is heavy. The objection concludes that, since the first disjunct is false, the latter must be true. The atomist response says that the objection ignores the possibility of plural property possession. The table may be identical with some atoms which are plurally heavy. This would allow for the claim that the table is heavy even if it is the case neither that each of the atoms is heavy, nor that the atoms constitute a complex object which is heavy.

Premise 2 is similar to the assumption made by the objection to atomism. It says that some things are collectively responsible for an action just if either each of them is responsible for it, or they constitute some complex entity which is responsible for it. As such, it is open to a challenge similar to that made by the atomist. Perhaps some things could be plurally responsible for an action. That is, perhaps we can deny the thought that responsibility for an action is necessarily a one-to-one relation between an action and an agent. Maybe responsibility is sometimes a one-to-many relation between an action and several agents. If this is true then some agents may be collectively responsible for an action even if it is neither the case that each of them is responsible for it, nor that they constitute a complex entity which is responsible for it.

²⁵ With the exception of Pinkert (2014, p. 189).

2.3. Conclusion

In this chapter, I have argued that both collective responsibility as several responsibility and collective responsibility as corporate responsibility may be compatible with responsibility as self-governance. Whilst (SR1) may not be consistent with claim (2), both (SR2) and (SR3) are consistent with (1) and (2). Moreover, given certain theories of collective agency, (CR) is also consistent with (1) and (2). As such, a deflationary response to the problem of collective responsibility is possible on both notions. Both several and corporate responsibility turn out to be ontologically committing; (SR3) minimally so, being committed plural actions, and (CR) strongly so, being committed to collective agents. At the end of Section 2.2.3, I hinted at the possibility of a further notion of collective responsibility: on which responsibility for some action(s) is possessed plurally by some agents. The following two chapters will discuss this notion.

Plural Responsibility

In Chapter 2 I discussed two notions of collective responsibility; collective responsibility as several responsibility and collective responsibility as corporate responsibility. In this chapter I will discuss a third possible notion of collective responsibility. I call this *collective responsibility as plural responsibility*. On this notion, several agents are collectively responsible for an action just if they are open to an instance of moral appraisal on the basis of the action(s). Here responsibility is treated as a property possessed plurally by several agents, following Section 1.4.2's (P).²⁶

In Section 3.1 I argue that collective responsibility as plural responsibility meets some necessary conditions for being consistent with responsibility as self-governance, and outline how it might meet some sufficient conditions. Section 3.2 provides a theoretical framework which, I argue, ensures that collective responsibility is consistent with responsibility as self-governance. Finally in Section 3.3 I discuss a problem which arises for Section 3.2's proposal. In this chapter I do not present examples in which several agents might bear collective responsibility as plural responsibility. Discussion of such cases is reserved for Chapter 4.

3.1. Plural responsibility and self-governance

Given the notion of collective responsibility as plural responsibility, we can clarify the possible responses to Section 1.3's problem of collective responsibility. To recap, the problem asks whether the notion of collective responsibility is consistent with the account of responsibility as self-governance, sketched in Section 1.2. On collective responsibility as plural

²⁶ A few authors have considered the possibility of collective responsibility of this form. Feinberg (1968, pp. 687-688) gives a brief discussion, whilst Jackson (1987) seems to have something similar in mind. The most extended discussion of something like collective responsibility as plural responsibility which I am aware of is in (Smith 2009, pp. 40-49).

responsibility, if some agents are collectively responsible for some action(s) then they are open to a single instance of moral appraisal. A radical response to the problem of collective responsibility says that responsibility as self-governance cannot give the result that several agents are open to an instance of moral appraisal. If some agents are open to an instance of moral appraisal on the basis of some action(s), then the action(s) cannot reflect their self-governance. As such, given collective responsibility as plural responsibility, collective responsibility is inconsistent with responsibility as self-governance. By contrast, a deflationary response says that if some agents are open to an instance of moral appraisal on the basis of some action(s) then the action(s) may reflect their self-governance. As such, collective responsibility is consistent with responsibility as self-governance.

In the previous chapter I argued that both several and corporate responsibility can allow for deflationary responses to the problem of collective responsibility. In this chapter, I will argue that the same is true for plural responsibility.

3.1.1. Agents and actions

In Section 1.3 I mentioned two features of responsibility as self-governance which authors have worried might conflict with features of collective responsibility. These were;

(1) Something is responsible for an action only if it is an agent.

(2) An agent is responsible for an action only if she performed it.

Neither of these claims, I suggest, conflict with collective responsibility as plural responsibility. We can further clarify this notion of collective responsibility as follows;

(PR) For every plurality of agents *xx*, *xx* are collectively responsible for some action(s) *aa* if and only if there is an instance of responsibility for *aa* which is possessed by *xx*.

Here responsibility is construed as a relation between some action(s) and some agent(s). When some agent(s) possess(es) an instance of responsibility for some action(s), a relation of responsibility is instantiated by the agent(s) and action(s). On (PR), when some agents are collectively responsible for some action(s), a single relation of responsibility is instantiated by the agents and the action(s). As such, if this notion of collective responsibility is coherent, responsibility cannot always be represented as a function from some action(s) to an agent. Sometimes responsibility can be construed as a one-to-many, or a many-to-many, relation between some action(s) and some agents.

Given this, on (PR), when several agents are collectively responsible for some action(s), it does not follow that any one thing is responsible for any action(s). To see this more clearly, consider again Section 1.4.2's example of the pile of books. The books are plurally heavy; there is an instance of heaviness which is possessed by the books. It does not follow from this that any one thing is heavy. It doesn't follow that any of the books is heavy; the books might be non-distributively heavy. Moreover it doesn't follow that the books constitute some complex object which is heavy. We might speak of the books as forming something grammatically singular; 'a pile', 'a plurality', 'a set', and so on. However this does not indicate that they constitute anything ontologically singular which is heavy.²⁷ Grammatical number does not indicate metaphysical number; a pair of trousers are one thing (Smith 2005, pp. 80-81).

Similarly when several agents possess an instance of responsibility for some action(s), it does not follow that something is responsible for some action(s). It doesn't follow that any one of the agents is responsible for some action(s). Moreover it doesn't follow that the agents constitute some complex entity which is responsible. As a result, (PR) is consistent with both claims (1) and (2). In order for either (1) or (2) to be false, it must be the case that something is responsible for some action(s). However, on (PR), if several agents are

²⁷ A putative exception is the set made up of the books. This set is something ontologically singular. However, whilst the books are heavy, the set made up of the books is not; a set is not the sort of thing that can be heavy. For a similar discussion of entailment see (MacBride 2005, pp. 581-582). See (Hossack 2000) for an attempt to avoid positing the existence of sets using insights from plural logic.

collectively responsible for some action(s) it doesn't follow that any one thing is responsible for any action(s).

One might modify claims (1) and (2) as follows;

(1-plural) Some things are responsible for some action(s) only if each of them is an agent.

(2-plural) Some agents are responsible for some action(s) only if they performed the action(s).

These claims seem to be plausible analogues of (1) and (2). I suggest that they are also entailed by responsibility as self-governance. Some action(s) cannot reflect the self-governance of some things unless they both possess such a capacity and performed the action(s) in question. Is (PR) consistent with claims (1-plural) and (2-plural)? I see no reason to think not. If some things possess a single instance of responsibility for some action(s), there is no reason to suppose either that those things are necessarily non-agents or that they did not perform the action(s) in question.

3.1.2. Plural self-governance

I have argued that (PR) meets some necessary conditions for being consistent with responsibility as self-governance; it is consistent with claims (1) and (2), and their plural analogues. However this is not sufficient for (PR) to be consistent with responsibility as self-governance. In addition, it must be the case that some action(s) is/are capable of reflecting the self-governance of several agents in a way that leaves them open to a single instance of moral appraisal.

In Chapter 2 I discussed two ways in which some action(s) might reflect the self-governance of several agents. Some action(s) might be indicative of each agent's capacity for self-governance. When this is the case, each of those agents is open to moral appraisal on the basis of the action(s); that is, they are collectively responsible for the action(s) on either (SR2) or (SR3), discussed in

Section 2.1.2. Alternatively some action(s) might be indicative of a capacity for self-governance possessed by some collective agent made up of the agents. In this case, the action(s) is/are a basis for moral appraisal of the collective agent; that is, the agents are collectively responsible for the action(s) on Section 2.2.1's (CR).

Neither of the above gives the result that the agents in question are open to a single instance of moral appraisal on the basis of the action(s) which reflect(s) their self-governance: that is, that they are collectively responsible for the action(s) on (PR). Is there some further way in which some action(s) might reflect several agents' self-governance? I suggest that there might be a third possibility. Some action(s) might be indicative of a capacity for self-governance which is possessed plurally by several agents. In such a case, the capacity for self-governance reflected by the action(s) is not possessed by an agent. Rather, the reflected capacity is possessed by several agents. If this were the case it would follow neither that the action(s) reflect(s) some capacities for self-governance possessed severally by the agents, nor that the action(s) reflect(s) some capacity for self-governance possessed by some complex thing made up of the agents. As such, it would follow that the agents are collectively responsible for the action(s) on neither (SR2), (SR3), nor (CR). Rather it seems plausible that they might be collectively responsible on (PR); that is, they would be open to a single instance of moral appraisal on the basis of the action(s).

Is my third suggestion feasible? Is it possible for several agents to possess a capacity for self-governance plurally? In Section 1.2.2 I discussed two important features of a capacity for self-governance. Firstly, a capacity for self-governance involves the ability to reflect on what is recommended or demanded by the practical reasons present in a given situation. Secondly, the capacity involves the ability to satisfy those recommendations or demands by acting, or not acting. If either of these two abilities can be possessed plurally by some agents, then I suggest that a capacity for self-governance can be possessed plurally.

Could several agents possess a sensitivity to reasons plurally? Being sensitive to reasons seems to require the possession of mental states. As discussed in Section 1.2.1, I take practical reasons to be evaluative facts. Being aware that such facts hold in a given situation seems to require being in possession of certain perceptual capabilities, being capable of forming beliefs, and so on. So if several agents can possess a sensitivity to reasons plurally, then it seems they must be capable of possessing certain mental states plurally. I have no decisive objections to the possibility of plurally possessed mental states. However the assumption that this is possible, and indeed that it actually occurs, is controversial.²⁸ If we can avoid this assumption, I suggest that we do so.

Might some agents possess an ability to satisfy the recommendations or demands of reasons plurally? In Section 1.2.1 I suggested that we conceive of a practical reason as a fact which supports, or opposes, an agent performing an action. An agent satisfies the recommendation or demand of a reason just if she performs an action which that reason supports, or does not perform an action which that reason opposes. Given this, it is possible for several agents to satisfy some reasons; by performing some actions each of which is supported by a reason, or by not performing some actions each of which is opposed by a reason. However in such cases some agents satisfy some reasons severally; for every one of the agents, there is an instance of satisfying a reason which is her satisfying a reason. This is not a case of some agents satisfying some reason(s) plurally; there is no instance of satisfying a reason which is possessed by the agents. In order for this to be the case, I suggest that we need to revise Section 1.2.1's conception of practical reasons.

3.2. Plural reasons

It is fairly common, I think, for someone to think about what several agents have reason to do, where it does not follow that she is thinking about what any one of them has reason to do. As A. J. Julius (2013a, p. 369) puts it: "It seems to me that I sometimes think about what we should do ... Just as there are

²⁸ See e.g. (Smith 2011, pp. 233-234; van Inwagen 1990, p. 5; Velleman 1997, p. 30).

facts about what I should do or about what you have reason to do, so there are facts about what you and I should do and about what you and I have reason to do". I take one of Julius' points to be that the facts about what you and I have reason to do are not always determined by the facts about what I have reason to do and what you have reason to do. This is not simply an *ad hoc* claim posited in order to justify some further claims. Rather, it is an observation about a way in which one might think about how several agents should act. In this section I explore one way of theorising about this observation. I argue that my claims can help to make sense of the notion of plural self-governance, and thus plural responsibility.

3.2.1. What is a plural reason?

In Sections 1.4.2 and 1.4.3 I presented two ways in which several things might possess a property and act. I suggested that some things might possess a property either severally or plurally, and that some agents might act either severally or plurally. I suggest that a similar distinction can be made in the case of several agents having a reason to act. A practical reason might be represented as a relation between an agent and the performance of an action; a reason either supports or opposes an agent performing some action. One might think that this relation is one-to-one; a practical reason relates an agent to her performing an action. Indeed, this was the position suggested in Section 1.2.1. However this might not always be the case. Perhaps sometimes a reason relates several agents to the performance of an action, or some performances of some actions. Given this, I suggest we can distinguish two ways in which several agents might have a reason to act;

(SRN) For every xx , xx severally have a reason to act if and only if for every x that is one of xx , there is a reason which x has to act.

(PRN) For every xx , xx plurally have a reason to act if and only if there is a reason which xx have to act.

According to (SRN), when a plurality of agents severally have a reason to act, for every one of the plurality there is a reason which supports her performing

some action. By contrast, on (PRN), when several agents plurally have a reason to act, there is a reason which supports the agents performing some action(s). Such a reason cannot be represented as a one-to-one relation between an agent and the performance of an action. Rather, it must be construed either as a many-to-one relation between several agents and the performance of an action, or a many-to-many relation between several agents and several performances of some actions. Notice that in (PRN) ‘plurally’ modifies ‘have a reason’ rather than ‘to act’, and the same for ‘severally’ in (SRN). That is, on (PRN), when some agents plurally have a reason to act, it is not necessarily the case that they have a reason to act plurally; they might plurally have a reason to act severally. Similarly, on (SRN), some agents might severally have a reason to act plurally.²⁹

The kind of reasons discussed in Section 1.2.1 either recommend or demand of an agent that she perform, or not perform, some action. Call these *singular reasons*. (PRN) posits a new kind of practical reason. Reasons of this kind are facts which either recommend or demand of several agents that they perform, or not perform, some action(s). Call such reasons *plural reasons*. To clarify, there are facts which appear to make several recommendations or demands, each of an agent that she perform, or not perform, some action. To illustrate, recall Feinberg’s sunbathing example, discussed in Section 2.1.2, in which a swimmer struggles in plain view of several sunbathers. If any one of the sunbathers comes to the swimmer’s aid, then the swimmer will be saved without excessive risk to the rescuer. This fact, I suggest, demands of each sunbather that she help the swimmer (conditional on enough others not attempting to help the swimmer). As such, there seems to be a reason which supports several agents performing some action(s). However such a reason is not a plural reason, on my use of the term. The fact that if any one of the sunbathers helps the swimmer then he will be saved makes several demands,

²⁹ Writers such as Alexander Dietz (2016), A. J. Julius (2013b, pp. 97-101), Robert Sugden (1993, pp. 84-85; Gold and Sugden 2007) and Thomas Smith (2009, p. 47) seem to discuss notions similar to (PRN). However, *contra* my own position, Dietz and Sugden appear to hold that if some agents plurally have a reason to act, then they have a reason to act plurally. Whilst he does not make this assumption, Julius (2013b, pp. 98-100), unlike myself, appears to hold that reasons had plurally can be reduced to reasons had severally.

each of a single agent. As such, the sunbathers severally have a reason to help the swimmer; for every one of the sunbathers, there is a reason which supports her helping the swimmer. By contrast, a plural reason is a fact which makes one recommendation or demand of several agents that they perform some action(s). Only when such a fact obtains will it be the case, on (PRN), that some agents plurally have a reason to act.

One might think that if some agents plurally have a reason to act, then it follows that each of them has a reason to ‘do their part’ of the action(s) supported by the plural reason.³⁰ This claim is equivalent to the following;

(α) For every *xx*, if *xx* plurally have a reason to perform some action(s) *aa*, then *xx* severally have a reason to perform *aa*.

I have no decisive considerations one way or the other as to whether this claim is true or false. However it should be noted that similar claims in the cases of property instantiation and action are false. If some books are plurally heavy, it does not follow that they are severally heavy. Similarly, if some musicians are playing plurally, it does not follow that they are playing severally: the score might indicate that some of the musicians rest. Given this, additional argumentation is required in order to establish that a similar claim for reasons is true. I do not have such an argument, nor do I have a decisive objection to the possibility of such an argument. However in Chapter 4 I discuss some cases in which it seems that some agents plurally, but not severally, have a reason to act. Such cases might pose a difficulty for defenders of (α).

3.2.2. Plural reasons and self-governance

On the assumption that some agents may plurally have a reason to act, I suggest that several agents can plurally possess a capacity for self-governance. As discussed above, such a capacity involves two important abilities; a sensitivity to reasons and an ability to satisfy reasons. I suggested that the

³⁰ For discussion see e.g. (Dietz 2016, pp. 968-982; Gold and Sugden 2007, pp. 125-128; Nefsky 2015).

former ability can be plurally possessed by some agents only given the controversial assumption that mental states can be plurally possessed. However, given that some agents can plurally have a reason to act, I suggest that they may plurally possess an ability to satisfy reasons.

A singular reason either recommends or demands of an agent that she perform, or not perform, some action(s). As such, a singular reason is satisfied just if an agent acts, or does not act, as the reason recommends or demands. By contrast, a plural reason either recommends or demands of several agents that they perform some action(s). As such, a plural reason is satisfied just if several agents act, or do not act, as it recommends or demands. Given this, an agent cannot satisfy a plural reason. If an agent performs some action(s) it does not follow that a plural reason has been satisfied, even if the action(s) is/are among those supported by a plural reason. Given the above, when several agents satisfy a plural reason it does not follow that any one of them has satisfied a reason; no one of them can satisfy a plural reason. As such, on (PRN), if some agents can plurally have a reason to act, it follows that they can plurally satisfy a reason.

To illustrate, suppose that some agents xx plurally have a reason r to perform some actions aa , and that xx are capable of performing aa . Given this, xx are capable of satisfying a plural reason; specifically, by performing aa . If xx satisfy r by performing aa , it will not follow that, for some x that is one of xx , x satisfies a reason. Even if, for some a that is one of aa , x performs a , x will not have satisfied r . This is because this reason either recommends or demands of xx that they perform aa . It demands nothing of any one of xx ; only a singular reason will recommend or demand something of an agent. Given this, it follows that xx are capable of satisfying a reason plurally. If they perform aa , as r supports them doing, they will have satisfied this reason plurally.

Of course if Section 3.2.1's (α) is true, then when xx satisfy r by performing aa it follows that each of xx satisfies a reason. Nevertheless it does not follow that if (α) is true then xx do not plurally satisfy a reason by performing aa . When xx satisfy r by doing aa , no one of xx satisfies r . This is the case

regardless of whether (α) is true. As such, even if (α) is true, when xx do aa there is an instance of satisfying a reason which is executed by no one of xx . Rather, it is executed by xx . Given this, even if (α) is true, xx plurally satisfy a reason by doing aa . The truth of (α) simply makes it the case that xx also severally satisfy some reasons by doing aa .

Given this, we have a clearer picture of the conditions under which several agents plurally possess a capacity for self-governance. Under these conditions, the agents are severally sensitive to reasons: singular and plural. That is, each agent is capable of recognising what is recommended or demanded by singular reasons of her, and others, in a given situation. She is also capable of recognising what is recommended or demanded of the agents by plural reasons. Furthermore the agents are capable of plurally satisfying any reason they might have plurally to act, or not act. That is, the agents are capable of satisfying the recommendation or demand of a plural reason, whilst no one of them is capable of doing so. Some agents exercise a capacity for self-governance which they possess plurally when they act, or don't act, in accordance with the recommendation or demand of a plural reason of which they are severally aware.

3.3. Reasoning about plurals

Reasons are things with which we reason. As J. J. Thomson (2001, p. 24) puts it: "A reason is something one might reason *from* [her emphasis]". A reason for a belief is something which can enter as a premise into some reasoning process, the end product of which is a belief. Similarly, a reason for action is something which one might input as a premise into some reasoning process, the end product of which is an action. Given the conceptual tie between reasons and reasoning, my account of plural reasons would be incomplete without a discussion of how one reasons with such reasons. In this section I discuss a problem which seems to arise when one inputs a plural reason as a premise in an instance of practical reasoning.

3.3.1. Sellars' Problem

In Section 3.2.1 I claimed that a plural reason recommends or demands of several agents that they act, or not act. It seems unproblematic that an agent might be sensitive to such reasons. I might recognise that some feature of a situation recommends or demands of some agent that she act, or not act.

Similarly, it seems plausible that I might recognise that some fact recommends or demands of some agents that they act, or not act. It also seems plausible that, on the basis of her sensitivity to plural reasons, an agent might reason to the conclusion that some agents should act, or not act. For instance, I might recognise that it is demanded of some agents they perform some actions *aa* and that it is not demanded of them that they act otherwise. I might therefore arrive at the conclusion that the agents in question should do *aa*.

Whilst both of the above claims seem plausible, they seem to lead to a problematic position. One might think that, just as the output of an instance of theoretical reasoning is a belief, the output of an instance of deliberation is an intention. As Natalie Gold and Robert Sugden (2007, pp. 110-111) put it: "When an agent deliberates about what she ought to do, the result of her reasoning is an intention. An intention is interposed between reasoning and action". Given this, it seems plausible to think that when I come to the conclusion that some agents should do *aa*, the output of my deliberation is the intention that the agents do *aa*.

However one might question whether the above intention is coherent. It seems plausible to think that an agent can intend something only if she believes that she can bring it about. That is, an agent cannot consistently both intend that *p* and believe that she can't make it that case that *p*. For instance, it seems inconsistent for me both to intend that Manchester United win their next game, and to believe that I can't bring it about that Manchester United win their next game. If I have this belief then I might *desire*, or *wish* for Manchester United to win, but I seem unable to consistently have this as an *intention*.³¹ To borrow

³¹ See (Sellars 1980, p. 83). Sellars (1980, p. 88) allows that a proposition may enter the content of an agent's intention by virtue of her believing that it will be true. However he maintains that "every intention must have at least one 'up-to-me' constituent" (1980, p. 89).

a phrase from Wilfrid Sellars (1980), I must view the content of my intentions as ‘up-to-me’. However an agent does not tend to believe that what others do is up-to-her. Given this, an intention of the form ‘I intend that several agents perform some action(s)’ seems to be incoherent. That is, I cannot consistently both intend that several agents act and believe that I cannot bring it about that they act. This problem was originally posed by Sellars (1980, p. 98). As such, I will refer to it as *Sellars’ Problem*.

3.3.2. Some solutions

How might this problem be avoided? One might deny that the output of every instance of practical reasoning is an intention. To illustrate, it seems plausible that I might be sensitive to reasons which you have to act. Furthermore I seem capable of reasoning, on the basis of this sensitivity, to the conclusion that you should perform some action *a*. But it seems plausible to think that the end-point of my reasoning is not an intention. As a result of my reasoning, I might wish or desire that you do *a*. It doesn’t seem unsatisfactory to think that I don’t form the intention that you do *a*.³² Given this, one might question why we should think that, when I reason to the conclusion that several agents should act, the end-point of my reasoning is the intention that several agents act. Perhaps it is satisfactory that my reasoning result in the wish or desire that several agents act. If so then Sellars’ Problem can be avoided. My having a wish or a desire that some agents act is consistent with my believing that I cannot bring it about that the agents act.

I allow that the output of some instances of practical reasoning might be a wish or a desire, rather than an intention. However I deny that this concession provides a satisfactory solution to Sellars’ Problem. There seems to be an important distinction between cases in which an agent reasons about what some agent should do, and cases in which an agent reasons about what she herself should do. To put it another way, there is a difference between cases in which I reason about what ‘someone’ should do, and cases in which I reason

³² I might intend to advise you to do *a*, or to help you to do *a*, or to manipulate you into doing *a*, and so on. However such intentions seem to result from my reasoning to the conclusion that I should advise you to do *a*, or that I should help you to do *a*, or ...

about what ‘I’ should do. In the former kind of case, it seems satisfactory that the end-point of my reasoning might be a wish or a desire. By contrast, in the latter kind of case this seems unsatisfactory. When I reason to the conclusion that I should act, I don’t merely form a wish or desire that I act. Rather, I form the intention to act. Indeed, this point seems implicit in the passage from Gold and Sugden (2007) quoted above.

Perhaps a similar distinction holds in the plural case. An agent might reason about what some agents should do. She might also reason about what some agents should do and include herself among the agents she is thinking about. That is, I might reason either about what ‘some agents’ should do or about what ‘we’ should do. Like in the singular case, there seems to be an important difference between these cases. Perhaps this distinction should be marked in a similar way. That is, when I reason to the conclusion that ‘some agents’ should act, the output of my reasoning may be a wish or desire that the agents act. By contrast, when I reason to the conclusion that ‘we’ should act perhaps I form the intention that we act. This proposal faces Sellars’ Problem; how can I coherently intend that we act?³³ However if the proposal is rejected, we must find another way of distinguishing cases in which I reason about what ‘we’ should do from cases in which I reason about what ‘some agents’ should do.

Distinguishing these forms of practical reasoning seems even more urgent given my interest in responsibility and self-governance. When an agent exercises her capacity for self-governance, she reasons about what she herself should do: not merely about what ‘some agent’ should do. Similarly if several agents exercise a capacity for self-governance which they possess plurally, it seems implausible that each will reason merely about what ‘some agents’ should do. Rather, each of them will reason about what ‘we’ should do; where ‘we’ refers to the agents who possess the capacity for self-governance in question.

³³ There is a growing literature on ‘we-intentions’. I will not review such accounts here. For discussion see e.g. (Bratman 1992; Gold and Sugden 2007; Kutz 2000, chap. 3; Miller and Tuomela 1988; Searle 1990; Smith 2011; Velleman 1997).

An alternative solution to Sellars' Problem might be as follows. Perhaps when an agent reasons about what 'we' should do, the result of her reasoning is an intention and a wish or desire. More specifically, when I reason to the conclusion that 'we' should do some actions *aa*, the result of my reasoning is my intention that I do *a* together with my wish or desire that others do *a'*: where *a* is/are the action(s) among *aa* which I believe I can perform and *a'* is/are the action(s) among *aa* which I do not believe I can perform.

This proposal allows us to distinguish cases in which an agent reasons about what 'some agents' should do from cases in which an agent reasons about what 'we' should do. In the former cases, the result of reasoning is merely a wish/desire that 'some agents' act, or not act. In the latter cases, the end-point of reasoning is a wish/desire about what others do plus an intention that the agent herself act, or not act. The proposal also avoids Sellars' Problem. It does not require that I intend anything which I do not view as up-to-me.

Specifically, it does not require that I intend that others act. Rather, I intend that I perform some action(s) which I view as up-to-me and wish or desire that others act.³⁴

However this suggestion faces a difficulty of its own. This difficulty becomes clear if we compare cases in which I reason about what 'we' severally have reason to do, and cases in which I reason about what 'we' plurally have reason to do. According to the above proposal, the results of both forms of practical reasoning will be the same. To illustrate, suppose I judge that 'we' severally have a reason to perform some actions *aa*: where 'we' refers to you and me. Given this, I may judge that I have a reason to perform some action *a* among *aa*, and that you have a reason to perform some other action *a'* among *aa*. Suppose that, on this basis, I judge that I should do *a* and that you should do *a'*. If so, then I will form the intention to do *a*. I may also form a wish or desire for you to do *a'*. That is, as a result of my reasoning, I intend that I do *a* and wish/desire that you to *a'*. Now suppose that I judge that 'we' plurally have a reason to do *aa*. Suppose further that, on this basis, I judge that we should do

³⁴ This proposal is similar to that suggested by Sellars himself (1980, p. 98).

aa. On the solution to Sellars' Problem proposed above, it follows from this that I will form the intention to do some action *a* among *aa*, and form a wish or desire that you do some other action *a'* among *aa*. We can clarify these two reasoning processes as follows;

Premise 1: We (you and I) severally have a reason to do *aa*.

Premise 2: From Premise 1 and (SRN) – I have a reason to do *a*, for some *a* that is one of *aa*, and you have a reason to do *a'*, for some *a'* that is one of *aa*.

Premise 3: From Premise 2 (and some other considerations) – I should do *a* and you should do *a'*.

Result: From Premise 3 – I intend that I do *a* and I wish/desire that you do *a'*.

Premise 1*: We (you and I) plurally have a reason to do *aa*.

Premise 2*: From Premise 1* (and some further considerations) – we should do *aa*.

Result*: From Premise 2* – I intend that I do *a*, for some *a* that is one of *aa*, and I wish/desire that you do *a'*, for some *a'* that is one of *aa*.

The results of these two reasoning processes, Result and Result*, are identical. Result follows from Premise 3 given the stipulation, made above, that if I judge that I should do *a* then I will form the intention to do *a* (and that if I judge that you should do *a* then I will form a wish/desire that you do *a*). However this stipulation does not make it the case that Result* follows from Premise 2*: Premise 2* says nothing about what I should do or about what you should do. Result* follows from Premise 2* only given the following stipulation; if I judge that we should do *aa* then I will form the intention that I do some *a* that is one of *aa* (and form the wish that you do some *a'* that is one of *aa*). But is this claim plausible?

It is plausible only if there are no cases in which I think both that we should perform some actions and that I should not perform any one of those actions. At least at first glance, such a case is not incoherent. Such a case would be one in which I think both that there is a plural reason which recommends or

demands of us that we do some actions *aa*, and that there are some singular reasons which recommend or demand of me that I not do any of *aa*. Given the above stipulation, in such a case I will form the intention to perform some of *aa*. I will also form the intention not to do any of *aa*. However the latter intention contradicts the former; I will intend both that I do some of *aa* and not do any of *aa*. It might be that, in actual fact, there are no cases like this. However in Section 4.2 I discuss some cases which might be described as such.³⁵

Notice that if Section 3.2.1's (α) is true then Sellars' Problem is avoided. If (α) is true then Premise 1* entails Premise 1. Given this, there is no problem of specifying what the result of Premise 2* is, since the premise can be by-passed in practical reasoning. As such, a potential solution to Sellars' Problem might be to argue for the truth of (α). As discussed in 3.2.1, I have no decisive considerations as to whether (α) is true. However there are cases, to be discussed in Chapter 4, which seem to be problematic for the claim.

I have reviewed two potential solutions to Sellars' Problem and argued that both face difficulties. Should we therefore discard the notion of a plural reason? I suggest such a response would be premature. Thinking about what 'we' should do seems to be a genuine mode of practical thought. Sellars' Problem should not be taken to demonstrate that this appearance is illusionary; it was certainly not intended as such by Sellars himself. Rather, I suggest that the problem should be taken as a puzzle concerning how to theorise about the plural mode of practical reasoning. I do not know how to respond to this puzzle. But this is no reason to think that it can't be responded to.³⁶

3.4. Conclusion

In this chapter I have argued that collective responsibility as plural responsibility can provide a deflationary response to the problem of collective

³⁵ See also (Jackson 1987; Smith 2009; 2011).

³⁶ A promising suggestion might be adapted from Smith (2011, pp. 235-242). Smith proposes that my (problematic) intention that 'we' (you and I) do some actions *aa* is distinct from my (unproblematic) intention that I do *aa* with you. Given this, perhaps what follows from Premise 2* is not Result* but Result**: I intend that I do *aa* with you.

responsibility. This notion of collective responsibility is consistent with claims (1) and (2), and their plural analogues. Furthermore I have argued that it is possible for some action(s) to reflect some agents' capacity for self-governance in such a way that leaves them open to a single instance of moral appraisal. This is the case when some action(s) reflect(s) a capacity for self-governance possessed plurally by some agents. When some agents exercise a capacity for self-governance which they possess plurally, they act, or do not act, in accordance with plural reasons. In exercising this capacity, the agents are severally aware of the reasons they have plurally to act, or not act. On the basis of this sensitivity, each deliberates about what 'we' should do. Some agents are open to an instance of moral appraisal on the basis of some action(s) just if the action(s) indicate(s) that the agents have, or have not, exercised their plural capacity of self-governance in accordance with the demands of plural, conclusive reasons.

Cases of Plural Responsibility

In Chapter 3 I introduced the notions of collective responsibility as plural responsibility and plural reasons. In this chapter I discuss some cases in which these phenomena might occur. This discussion is intended to strengthen and clarify the claims of the previous chapter. I focus on puzzle cases in which, whilst it seems plausible that the agents involved are open to moral criticism on the basis of their actions, it also seems plausible that no agent has a reason to act otherwise. These two claims seem to contradict each other. However if we assume that the agents in such cases plurally have a reason to act otherwise, then the two plausible-sounding claims can be retained without contradiction. In Section 4.1 I discuss a class of cases I call ‘coordination problems’, before examining a case originally discussed by Frank Jackson (1987) in Section 4.2. Section 4.3 discusses a class of cases I call ‘no difference cases’. Section 4.4 concludes with a brief discussion of consequentialism.

4.1. Coordination problems

In this section, I discuss cases in which it is open to several agents to perform severally some possible combinations of actions. The agents face a choice as to which of the combinations to perform. Some of these combinations appear more advisable than others. As such, the thing for the agents to do appears to be to coordinate their actions such that they perform one of the more advisable combinations. However it also seems as though no one of the agents has a reason to perform their part of one the advisable-looking combinations. As such, the thought that some combinations of actions are more advisable than

others turns out to be difficult to justify. Call such cases *coordination problems*.³⁷

Coordination problems occur frequently in the actual world and tend to be successfully solved by the agents involved; that is, when faced with such cases, real agents tend to choose the more advisable-looking action-combinations. As such, the conclusion that agents in coordination problems have no reason to choose the advisable-looking combinations over the others seems puzzling. I will argue that the problematic appearance of such cases is resolved if we assume that the agents plurally have reason to perform the advisable-looking combinations. My discussion of coordination problems broadly follows that of Julius (2013b, chap. 10).

To illustrate consider the following example;

The Ambulances: Two ambulances are driving in opposite directions down a narrow road, when they come face-to-face. One ambulance driver (call her ‘Driver N’) is taking an elderly woman who has suffered a heart attack to a specialist cardiac intensive care unit to the north, whilst the other (call her ‘Driver S’) is carrying someone who has had a stroke to a specialist brain injury unit in the south. Both individuals will die unless they are treated in their respective care units soon. Since the road is so narrow, the ambulances can pass each other only if one ambulance driver steers in one direction whilst the other steers in the opposite direction. There are no laws or conventions governing which way the ambulances should steer. As it turns out, neither driver steers in either direction. The ambulances sit facing each other and both patients die.

³⁷ Coordination problems have been widely discussed in game theory. In game theory, a coordination problem is represented as a game in which several players face a choice between several combinations of their singular actions. Some of these combinations have higher utility payoffs for each agent than others. However, in standard game theory, the claim that the agents should coordinate their actions so as to achieve a higher utility payoff cannot be justified. See Sugden (1993) for further discussion.

The ambulance drivers' decision(s) strike(s) us as a failure. Moreover the decision(s) seem(s) to strike us as a *moral* failure. Had the drivers chosen to steer around each other, the patients' lives could have been saved. That the drivers did not do this does not seem merely foolish or unwise. Rather, it looks like a basis for moral criticism.³⁸ However I suggest that this intuitive thought cannot be justified by appeal to singular reasons.

Given Section 3.2.1's (SRN), the drivers severally have a reason to steer around each other just if Driver N has a reason to steer to one way and the Driver S has a reason to steer the other way. What reason might either driver have to steer in either direction? It seems any reason a driver has to steer one way or the other will be conditional on what the other does. For instance, Driver N has a reason to steer west just if Driver S steers east; when, and only when, Driver S steers east, Driver N can continue northward by steering west. As such, neither driver has a reason to steer one way or the other unless the other driver steers in the opposite direction. But if this is correct then neither driver fails to do what she has reason to do when, as it turns out, both of them stay still. Since Driver S moves neither east nor west, Driver N has no reason to steer in either direction: and *vice versa*.³⁹

One might make the following objection. Suppose Driver N has a reason to steer west if Driver S steers east. Suppose further Driver S has a reason to steer east if Driver N steers west. Finally, suppose that both drivers will do whatever they have reason to do. Given this, it seems to follow that if Driver N steers west then Driver S will steer east. As such, if Driver N steers west then she will have a reason to steer west. This in itself, one might think, is a reason for Driver N to steer west. This argument can be presented more formally as follows;

Premise 1: Driver N has a reason to steer west if Driver S steers east.

Premise 2: Driver S has a reason to steer east if Driver N steers west.

³⁸ By contrast, 'unsuccessful' responses to standard game-theoretic coordination problems do seem merely foolish. See (Pink 2007, p. 425).

³⁹ See Julius (2013b, p. 94).

Premise 3: Driver S will do whatever she has reason to do.

Premise 4: From Premises 2 and 3 – if Driver N steers west then Driver S will steer east.

Premise 5: From Premises 1 and 4 – if Driver N steers west then Driver N will have a reason to steer west.

Premise 6: If by ϕ -ing an agent will make it the case that she has a reason to ϕ , then she has a reason to ϕ .

Conclusion: From Premises 5 and 6 – Driver N has a reason to steer west.

If this argument, and analogues for Driver N steering east and Driver S and steering in either direction, succeeds then we can justify the thought that the drivers are open to moral criticism. The drivers severally have a conclusive reason to steer around one another. When they both stay still, they thereby fail to perform some actions that are demanded of them.

I suggest that this proposal fails. Julius (2013b, p. 95) worries that, if the argument succeeds, Driver N's action may be impermissible. If Driver N steers west with the aim of making Driver S steer east, whilst simultaneously making it the case that Driver S has a reason to steer east, then Driver N's action may count as coercive.⁴⁰ In what follows I raise a different concern about the argument's validity.

The above argument supposes that Premise 4 is entailed by Premises 2 and 3. That is, given that Driver N steering west provides Driver S with a reason to steer east and that Driver S will do whatever she has reason to do, then if Driver N steers west then Driver S will steer east. However this does not follow. It follows only on the assumption that Driver S does not have a reason to steer west. If she does then, even assuming Premises 2 and 3, Premise 4 will not be true. Given Premise 3, and that Driver S has a reason to steer both west and east, then Driver S might steer either east or west.

⁴⁰ For Julius' account of coercion see (Julius 2013a).

However, assuming that the above argument is sound, it seems plausible that Driver S does have a reason to steer west. This result can be derived by substituting in the above argument 'Driver N' for 'Driver S' and *vice versa*. A proponent of the above argument thereby faces a problem. It cannot be the case that both the above argument and its substitution instance for Driver S are sound. So one may hold that the above argument is sound only if one denies that the substitution instance is sound. But this seems like an *ad hoc* position. What could make it the case that the above argument succeeds whilst its substitution instance for Driver S fails? A more plausible position seems to be to deny that either argument is sound.⁴¹

I have argued that it is not the case that the ambulance drivers severally have a reason to steer around each other. Given this, it follows that the drivers are not collectively responsible for their inaction on collective responsibility as several responsibility. That is, it is not the case that each driver is open to moral criticism on the basis of their inaction. Moreover it seems unlikely that the drivers will be collectively responsible on collective responsibility as corporate responsibility. I suggest that it is not the case that the drivers constitute some complex entity which is open to moral criticism on the basis of its inaction. One might specify that there is some complex entity made up of the drivers. However it seems unlikely that such an entity will possess a capacity for self-governance. As such, the claim that this entity is attributably responsible for some action(s) or inaction(s) will conflict with responsibility as self-governance. How then can we justify the thought that the drivers are open to moral criticism?

I suggest that this claim can be justified if we suppose that they plurally have a conclusive reason to steer around each other. I propose that there is a fact which demands of the drivers that they perform severally some plurality of actions in which one driver steers in one direction whilst the other steers in the

⁴¹ Sugden (1993, pp. 74-75) argues that coordination problems pose a similar problem for standard game theory. He claims that, for a player in a coordination problem, an answer to the question 'what should I do?' will be viciously circular: since what I should do is conditional on what the other player should do.

opposite direction. Such a fact might be along the following lines; if the drivers steer around each other, then the patients' lives will be saved. This fact does not demand of either driver that she steer in any direction. Rather, it demands of the drivers plurally that one of them steers one way and the other steers the other way.

Given this, we can say that the ambulance drivers are collectively responsible for their inaction on (PR). That is, the drivers' inaction is a basis for moral appraisal of them plurally. If the drivers plurally have a conclusive reason to drive around each other, then by staying still they fail to perform some actions which are demanded of them. This failure to act reflects neither driver's capacity for self-governance. Rather, their inaction reflects a capacity for self-governance which the drivers possess plurally. As a result they are plurally open to moral appraisal on the basis of their inaction; that is, on (PR), they are collectively responsible for their inaction.

4.2. Jackson's case

In this section, I will discuss cases with a similar structure to coordination problems. In such cases, some agents must choose between several possible combinations of actions. As in coordination problems, some of these combinations seem more advisable than others. Moreover, unlike coordination problems, one of these combinations seems most advisable. But it also seems as though it's not the case that the agents severally have a reason to perform the most advisable-looking combination. Given this, the thought that the agents should choose the most advisable-looking combination is hard to justify. I suggest that this problem is resolved if we assume that agents in such cases plurally have a reason to perform the most advisable-looking combination. My discussion in this section broadly follows that of Smith (2009).⁴²

⁴² These cases are structurally similar to Hi-Lo games discussed in game theory. In a Hi-Lo game, several players must choose between several combinations of actions. One of these combinations has a higher utility payoff for each agent than the others. However, in standard game theory, the result that the agents should choose the highest paying combination cannot be proved. The cases discussed in this section are also similar to Prisoner's Dilemmas. See (Gold and Sugden 2007; Smith 2009, pp. 38-39; Sugden 1993) for further discussion.

I will focus my discussion on the following case, originally from Jackson (1987) and adapted by Smith (2009);⁴³

The Drivers; 1000 individuals are driving their cars on a motorway. Their cars are tuned such that they can move only at either 60 mph or 80 mph. If all of the drivers drive their cars at 60 mph then the risk of a serious traffic accident will be relatively low. If they all drive at 80 mph then the risk of an accident will be significantly higher. However it is most dangerous for some to drive at 60 mph whilst some drive at 80 mph, due to the overtaking that will result. There are no traffic laws regulating the speed at which the drivers may drive. As such, no driver does anything illegal by driving at either 60 mph or 80 mph. Furthermore no one derives any significant benefit from any one of the drivers driving at 80 mph rather than 60 mph. No driver is in any particular rush to get to their destination, no driver gains any significant pleasure from driving at a higher speed, and so on. As it turns out, all of the drivers drive their cars at 80 mph thereby realising the second safest option.

It seems plausible to think that the drivers are open to moral criticism on the basis of their actions. It seems an imperative that the drivers keep the risk of a serious accident as low as possible. If they all drove at 60 mph then this imperative would be satisfied. However the drivers did not do this; they opted to each drive at 80 mph. By doing so, they increased the risk of an accident to a higher level than necessary.

However, due to considerations similar to those discussed in Section 4.1, it also seems as though it's not the case that the drivers severally have a conclusive reason to slow down. It seems as though, for any one of the drivers, that driver has a reason to drive at 60 mph just if all of the others also drive at 60 mph. Given that the others will drive at 60 mph, by also driving at 60 mph the driver can bring about the safest plurality of actions available. However if any of the others drive at 80 mph, then the driver cannot make the situation

⁴³ For discussion of similar cases see e.g. (Dietz 2016; Feldman 1980; Jackson 1988; Pinkert 2015; Regan 1980).

any safer by driving at 60 mph. Indeed when everyone else drives at 80 mph, the driver will make the situation significantly less safe by driving at 60 mph. As such, when everybody drives at 80 mph, it follows that no one has a reason to drive at 60 mph. Indeed it seems as though each driver has a conclusive reason not to drive at 60 mph.

Given this, it follows that the drivers cannot be collectively responsible for their actions on collective responsibility as several responsibility. Since no driver has a conclusive reason to slow down, it cannot be the case that the actions reflect each driver's capacity for self-governance. Furthermore it seems implausible that the drivers constitute some complex entity which is responsible for the actions. Even if they do constitute something, it seems plausible that there is at least one description of the case on which that thing does not possess a capacity for self-governance. As such, I suggest that the drivers are not collectively responsible for driving at 80 mph on collective responsibility as corporate responsibility.

However I suggest that the agents might be collectively responsible for their actions on collective responsibility as plural responsibility. I propose that the drivers plurally have a conclusive reason to each drive at 60 mph. That is, given (PRN), there is a fact which demands of the drivers that they severally drive at 60 mph. Such a fact may be something like the following; if the drivers severally drive at 60 mph, then the risk of a serious traffic accident will be minimised. This fact does not demand of any one of the drivers that she drive at 60 mph. Rather, the demand is made plurally of the drivers (Smith 2009, p. 47).

Given this, on (PR), the drivers are collectively responsible for severally driving at 80 mph. That is, the drivers are open to an instance of moral appraisal on the basis of their actions. This is because these actions reflective of a capacity for self-governance which the drivers possess plurally. Their severally driving at 80 mph indicates that they have not exercised this capacity in accordance with plural, conclusive reasons.

One might have the following worry. If the drivers plurally have a reason to drive at 60 mph, then this reason gets violated if all but one of the drivers drive at 60 mph. Since, in this case, the drivers have not acted in accordance with a plural, conclusive reason, one might think that the drivers are plurally open to moral criticism on the basis of those actions. However this result might seem implausible. Only one driver seems open to moral criticism in this case; the one who drove at 80 mph. The others who drove at 60 mph do not seem open to moral criticism.

However notice that in this case, the speeding driver has a reason to slow down; if she slows to 60 mph then she will bring about the safest outcome. As such, there are two conclusive reasons the demands of which are violated; the plural reason and the singular reason. The question is: which capacity for self-governance is reflected in this case? Is it the speeding driver's, or the capacity possessed plurally by the drivers? The case might be re-described in different ways such that either seems plausible. However the kind of case imagined by the above worry is, I suggest, one in which the speeding driver's self-governance is reflected by her action. In such a case the slow drivers, sensitive to the reason they have plurally and reasoning about what 'we' should do, decide to drive at 60 mph. However the speeding driver, either due to a lack of sensitivity to reasons or because of a conscious decision to violate to demand of this reason, fails to act in accordance with the reason she has to slow down. In such a case, the drivers are not collectively responsible for their actions on (PR). Rather, the speeding driver is responsible for her own action.

4.3. No difference cases

In this section I will discuss a class of cases which I call *no difference cases*. These are cases in which if a large number of agents each perform an action of a particular type, then a particularly objectionable outcome will be produced. However no one of these actions makes a difference to the outcome in question. For any one of these actions, when the action is performed, the outcome is precisely the same as it would have been had that action not been performed. No difference cases seem to be prevalent throughout the world today. For instance, anthropogenic climate change is caused by millions of

individuals performing actions which require the emission of greenhouse gasses. However it seems to make no difference to the effects of climate change whether or not I choose to drive my car to work today (Sinnott-Armstrong 2005). Again, the existence of sweat-shops is, in part, caused and perpetuated by millions of individuals demanding cheap clothing. Yet it seems to make no difference to sweat-shop conditions whether or not I choose to buy a pair of £5 jeans (Young 2004; 2011).⁴⁴

Whilst discussion of such real-world cases is important, for conceptual clarity I will focus on the following abstract case;

The Reservoir: 1000 individuals (the tippers) each has some waste they wish to dispose of. There are various places in which they could dump their waste. One of the most convenient places in which to do so is a reservoir. However there is a community who live nearby and the reservoir is their main source of water. If all (or enough) of the tippers dispose of their waste in the reservoir then the community will be poisoned. Moreover if all (or enough) of the tippers dump their waste elsewhere then the community will live in good health. But one tipper's disposing of her waste in the reservoir will make no difference to the community's health. For any number n from 1 to 1000, if n tippers dump their waste in the reservoir then the health of the community will be precisely the same as it would have been had $n-1$ tippers dumped their waste in the reservoir.⁴⁵ No tipper derives any significant benefit from disposing of her waste; they'd simply prefer to be rid of it. There is also no law which prohibits dumping waste in the reservoir. Furthermore dumping waste in the reservoir will have no adverse effect on any wildlife or on the natural beauty of the surroundings; the reservoir is devoid of wildlife and is quite ugly. As it turns out, each tipper disposes of her waste in the reservoir and the community is poisoned as a result.

⁴⁴ For further discussion of no difference cases see e.g. (Glover 1975; Kagan 2011; Nefsky 2011; Parfit 1984, chap. 3; Speikermann 2014).

⁴⁵ This appears to lead to a sorites paradox. See (Kagan 2011; Nefsky 2011; Speikermann 2014) for discussions of sorites paradoxes and vagueness in the context of no difference cases.

This case seems puzzling. On the one hand, it seems plausible to think that the tippers are open to moral criticism for severally dumping their waste in the reservoir, since the community were poisoned as a result of their actions. On the other hand, it also seems plausible that no tipper has a reason not to dispose of her waste in the reservoir. This is because one singular act of dumping waste in the reservoir makes no difference to the health of the community. As such, there is no fact along the lines of ‘if this action is performed then someone’s health will suffer’ which opposes a tipper dumping waste in the reservoir. Given this, it follows that no tipper is open to moral criticism on the basis of her action. I suggest that by introducing the notion of a plural reason, we can go some way towards resolving the problematic appearance of cases like this one.

4.3.1. No difference cases and singular reasons

Before showing how the notion of a plural reason might help to resolve the apparent problem with no difference cases, I will make a case for thinking that the agents in such cases do not severally have a conclusive reason to act. In what follows, I discuss four proposals which aim to establish that a tipper has a conclusive reason to abstain from dumping in the reservoir. Each of them, I argue, fails.⁴⁶

Probabilities: Suppose that I am a tipper. What reason might I have to abstain from dumping my waste in the reservoir? It’s clear that my reason for abstaining cannot be that if I abstain then I will prevent the community from being poisoned. One act of dumping makes no difference to the community’s level of health. So my abstaining from dumping cannot make the difference between the community being poisoned and the community enjoying a satisfactory level of health. However one might think that by abstaining from dumping, I *might* prevent the poisoning of the community. That is, my reason might be that if I abstain, my abstaining will have some probability of preventing the community from being poisoned.

⁴⁶ Julia Nefsky (2015) has given an extensive review of the reasons an individual might have to act in no difference cases, arguing that all the proposals she considers fail. In what follows I aim to supplement Nefsky’s discussion.

But this proposal can't be right. It assumes that one act of dumping can make the difference between the community being poisoned or not. Elections have a structure like this. In most elections with large numbers of voters, one vote makes no difference to the result. But it is possible that this might happen. This is so in the case of a tie or a one-vote-win. In such cases, every vote makes a difference; each vote is such that, had that vote been different or not been made, the outcome of the election would have been different. Given this, every vote in every election has some probability of making a difference: the probability of there being a tie or a one-vote-win. However I have assumed that no difference cases are not like this; whether or not I dump my waste in the reservoir could never make a difference to whether or not the community is poisoned.⁴⁷

Conditional reasons: Recall that in coordination problems and in cases like Jackson's drivers, an agent has a reason to act in a certain way just if others act in a certain way. Call such a reason a *conditional reason*. As discussed in Sections 4.1 and 4.3, introducing the notion of a conditional reason does not help to justify the thought that agents in such cases make a mistake when they choose to perform putatively inadvisable combinations of actions. However it does allow us to justify the thought that agents make a rational choice when they choose putatively advisable combinations of actions. For instance, if the ambulance drivers in Section 4.1's ambulance case drove around each other, we can justify the thought that they were right to have chosen that combination of actions using conditional reasons; since one driver steered her ambulance in one direction, the other driver had a reason to steer in opposite direction, and *vice versa*.

One might think that no difference cases are similar to coordination problems and Jackson's driving example in this respect. Perhaps I have a conclusive reason to abstain from dumping my waste in the reservoir which is conditional

⁴⁷ Shelly Kagan (2011) gives a response to no difference cases like the one I have just countered. He argues that all putative no difference cases in fact have a structure like that of elections. As such, he denies that there could be any no difference cases of the sort I have described. See (Nefsky 2011) for an argument that, *contra* Kagan, such cases are possible.

on enough others also abstaining. This suggestion will not help justify the thought that the tippers are open to moral criticism when they all dump their waste in the lake. But it can justify the claim that they get something right if enough of them abstain from dumping in the lake; and perhaps we can rest content at that.

However there is a problem with this proposal. This problem arises in trying to specify what reason there might be for me to abstain from dumping my waste in the reservoir which is conditional on the actions of others. In Section 4.1's ambulance case, if one ambulance driver steers to her left, then the other can deliver her patient to the hospital if she steers to her left. That is, the latter ambulance driver's reason for steering to her left is that she can save her patient's life if she does so. But she can't achieve this by steering to her left unless the other driver steers to her left. As such, her reason for steering to her left is conditional on the other driving steering to her left.

Perhaps my reason for abstaining from dumping in the reservoir is like this; if the other tippers abstain, then I can achieve something that I previously couldn't by abstaining. But what could this achievement be? I can't prevent the community from being poisoned by abstaining, regardless of how many other tippers also abstain. As already discussed, whether or not I dump my waste in the reservoir can never make a difference to the health of the community. But as such, it's unclear what I might achieve by abstaining which could give me a conditional reason to do so. It seems as though all I would achieve by abstaining would be to ensure that there is slightly less waste in the lake. But it's not clear why this would be a reason for me to abstain, conditional on others abstaining.

Fairness: One suggestion here might be that, if all (or enough) of the other tippers abstain from dumping their waste, then it would be unfair for me not to abstain as well. That is, my dumping in the reservoir would be unfair, conditional on others abstaining. Scanlon (1998, pp. 211-213) has made such a

suggestion.⁴⁸ Scanlon considers a cooperative venture which produces a certain good, but which is such that it makes no difference whether or not one member of the cooperative pulls her weight; if one person doesn't participate, the good will still be produced at the cost of no extra effort from the other participants. Nevertheless, Scanlon claims that a principle which required all members of the cooperative but one to participate in the venture would be unfair. This is because every member has reason to want a principle which requires everybody but themselves to participate to be implemented; the good of the venture will be produced and they won't have to do any work to get it. But a principle which calls for everyone but some individual *i* to participate favours *i*'s claim for exemption over everyone else's. So to implement this principle, without further justification, would be unfair; it would favour *i*'s claim for exemption over others without justification.

No difference cases seem to have a similar structure to Scanlon's cooperative scheme. If enough others abstain from dumping their waste in the reservoir, the community will enjoy a satisfactory level of health whether or not I dump my waste in the reservoir. Furthermore nothing more will be required of the others to prevent the community from being poisoned if I dump in the reservoir. However we can suppose that each tipper has reason to want it to be the case that all, or enough, tippers abstain, whilst she herself is exempted; the community won't get poisoned and she will be rid of her waste. So, following Scanlon, one might argue that if all, or enough, of the other tippers abstain from dumping in the reservoir, it would be unfair if I did not also abstain. In this case, if I dumped in the reservoir I would be favouring my own claim to be exempted from abstinence over those of others without further justification. Given this, there is a conditional reason for me to abstain from dumping my waste in the reservoir. It would be unfair if I didn't abstain, given that others will do so.

⁴⁸ Garrett Cullity (2000) has made a similar suggestion, to which Nefsky (2015) has responded. My response to Scanlon draws on Nefsky's discussion.

I suggest that this proposal fails. The argument relies on the claim that, by dumping in the reservoir whilst others abstain, I favour my own claim to be exempted from abstinence over others' *without justification*. However there does seem to be a justification available; specifically, that my abstaining will make no difference to the community's health. To see this, consider the following example from Nefsky (2015, p. 257). A car needs pushing up a hill. There are five volunteers willing to do the pushing, and their pushing together would be sufficient to get the car up the hill. However one of the five is extraordinarily weak. His pushing won't contribute anything to the joint effort; if he doesn't push, the other four won't have to push any harder to get the car up the hill. It does not seem unfair if the weak volunteer exempts himself from pushing the car. He has reason to want this exemption; the car will still get up the hill and he won't have to do any pushing. But he also has a justification for his exemption; his pushing would be instrumentally superfluous to the joint effort.

In the reservoir case, I seem to have a similar justification for exempting myself from abstaining from dumping waste in the reservoir. If enough others abstain, this will be sufficient to prevent the community from being poisoned. Furthermore, if I don't abstain with them, the others won't have to do anything more in order to prevent the community from being poisoned. Given this, it seems as though my abstaining will be instrumentally superfluous to saving the community from being poisoned.⁴⁹ If so, then I have a justification for exempting myself. It's true that every other tipper also has the same justification for exempting themselves. But this does not make the justification unsound. What follows from this is that for every pattern of actions in which enough tippers abstain in order to prevent the community from being poisoned, others have a justification for exempting themselves from abstaining. Any one of these patterns of actions can be implemented without the exemption(s) being unfair.

⁴⁹ Nefsky (n.d.) argues that we can distinguish between an action's not making a difference to some outcome, and its being instrumentally superfluous to bringing that outcome about. Whilst an action might not make a difference to whether an outcome occurs, it may not be instrumentally superfluous to this. I will not discuss Nefsky's proposal here.

Cooperative schemas: I want to discuss one final suggestion. There are various cooperative schemas such that if they were implemented, then I would have a reason to abstain from dumping my waste in the reservoir. For instance, suppose that each tipper agrees not to dump in the reservoir if and only if every other tipper does so too. Under this cooperative schema, call it *C*, each tipper's abstaining is necessary to prevent the community from being poisoned. So under *C* there is a reason for me to abstain: namely, if I don't abstain then the community will be poisoned. As I described it above, no such cooperative schema has been implemented in the reservoir case. As such, I have argued that it is very difficult to see what conclusive reason I might have to abstain from tipping my waste into the reservoir. However one might think that, whilst I have no reason to abstain in the absence of a cooperative schema such as *C*, there is a reason for me to do something else: specifically, to implement a cooperative schema like *C*. My reason for doing so would be something like the following; if I implemented *C*, then I would increase the probability of the community enjoying a satisfactory level of health: since, under *C*, everyone would have a conclusive reason not to dump their waste in the reservoir.

This proposal seems promising. However, as it stands, I don't think it is successful. To see this, consider how I would go about implementing a cooperative schema like *C*. How do I do this? Maybe I could do it by addressing the tippers and saying something like "OK everyone, what I propose it this ..." and proceed to describe *C*, or some similar schema. Maybe it includes things like persuading the tippers that the schema would be a good idea, or writing up contracts for the tippers to sign, and so on. These sorts of thing do seem to be involved in the implementation of the schema. But this can't be what it is to implement the schema. *C* gets implemented only when each tipper agrees not to dump in the reservoir if and only if every other tipper does so too. Similar claims go for similar schemas. Whilst I can try and make this happen, I can't actually do it. I cannot agree to the terms of the schema for you; only you can do that. What this reveals is that I cannot implement a cooperative schema like *C*. This is not an action which I can perform. Rather, it is either a plurality of actions performed severally by some agents, or a

plural action. Given this, I cannot have a reason to implement a cooperative schema like *C*. It cannot be the case that there is a fact that demands of me that some agents act; this is incoherent. A reason which makes a demand of me must demand that *I* do, or not do, something.

I can, however, do some part of the plurality of actions, or plural action that is the implementation of *C*. I can propose the schema to the tippers, I can persuade them to sign up to it, I can sign up myself, and so on. Maybe there is a reason for me to do these things. Notice that my reason for doing these things cannot be the reason I would have to implement the schema, if I could do such a thing. My reason to implement the schema would be that if I did so, I would increase the probability of the community not being poisoned. But my doing some part(s) of the implementation of the schema might not have this consequence. For instance, if I signed up to the schema but none of the other tippers agreed to do so too, then I would not have increased the probability of the community not being poisoned.

In light of this, one might think that I have a conditional reason to do some part(s) of the implementation of the schema. That is, perhaps there is a reason for me to sign up to the schema, for instance, just if the others sign up as well. This might well be correct. However, for reasons similar to those discussed in Sections 4.1 and 4.2, this proposal cannot justify the thought that the tippers are open to moral criticism when they don't implement a cooperative schema. If we want to account for this thought, we need something more than the claim that the tippers severally have a conditional reason to implement a cooperative schema like *C*.

4.3.2. No difference cases and plural reasons

I have argued that it is not the case that the tippers severally have a reason not to dump their waste in the reservoir. But perhaps they plurally have a reason not to do so. More specifically, perhaps there is a plural, conclusive reason which demands of the tippers that they not severally dump their waste in the reservoir. Such a reason might be a fact such as the following; if the tippers severally dump waste in the reservoir then the community will be poisoned.

This fact demands nothing of any single tipper. Rather, the demand is made plurally of the tippers. Alternatively, the tippers might plurally have a reason to implement a cooperative schema like *C*.

Given this, it is possible that the tippers might be open to moral criticism on the basis of their actions. By dumping in the reservoir, the tippers act contrary to the demand of a plural, conclusive reason. As such, their actions may indicate that they have failed to exercise as they should a capacity for self-governance which they possess plurally. If so then the tippers will be open to an instance of moral criticism on the basis of their actions. That is, the tippers will be collectively responsible for their actions on (PR).

This proposal might seem to be a promising solution to the puzzle surrounding no difference cases. However it faces a problem. If all of the tippers dump waste in the reservoir then the community will be poisoned. I suggested that this fact might be a plural reason which demands of all the tippers that they not dump waste in the reservoir. However the following fact also seems to hold; if all but one of the tippers dump waste in the reservoir then the community will be poisoned. Consequently, if all the tippers plurally have a reason not to dump in the reservoir, then it looks like all but one of the tippers will also plurally have a reason not to dump in the reservoir (notice that, since there are 1000 tippers, there are 1000 pluralities which satisfy the description ‘all but one of the tippers’). The same goes for all but two of the tippers, all but three of the tippers, and so on.

This presents a difficulty for the above proposal. There is a vast number of pluralities of tippers, each of whom it seems true that they plurally have a reason not to dump waste in the reservoir. Given this, why should we think that the tippers’ actions reflect a capacity for self-governance possessed plurally by all the tippers? Why not think, for example, that a capacity for self-governance possessed plurally by all but one of the tippers is reflected by some actions? As such, it is unclear why we should think that some plurality of tippers, rather than another, are collectively responsible for their actions on (PR).

I will not explore a response to this problem here. However I suggest that it is not a decisive objection to employing the notions of a plural reason and plural responsibility in discussions of no difference cases. These notions do not resolve the problematic appearance of such cases. However they may contribute towards such a solution. These notions, I suggest, help to clarify what the problem with no difference cases is. The problem is this; which, if any, plurality are responsible for their actions? Notions of plural reasons and responsibility do not solve this problem. But they help us to pose it.

4.4. Consequentialism

I have claimed that, in each kind of case discussed in this chapter, the agents involved appear to be open to moral criticism on the basis of their actions. However I have argued that no agent has a reason to act otherwise. Given this, no agent is open to moral criticism on the basis of her action(s). I have argued so by appealing to the consequences of their actions. As such, one might think that my claims reveal an implicit commitment to a specific moral theory; namely, consequentialism. This might be fine; various discussions of the sorts of cases discussed here are explicitly consequentialist.⁵⁰ But one might worry whether someone who has no commitment to consequentialism should be persuaded by my arguments.

This worry is premised on a misunderstanding of consequentialism.

Consequentialism is not the view that, in determining what an agent ought to do, we should appeal to the consequences of her actions. As David Wiggins (2006, p. 149) points out; “Every sane moral philosopher since the beginning of the subject has had regard for outcomes and consequences...If an ethical argument appeals to outcomes, that need not make it into a utilitarian or a consequentialist argument”. Rather, the term ‘consequentialism’ was introduced by Elizabeth Anscombe (1958) to pick out a very specific version of this uncontroversial view; “that ‘the right action’ is the action which produces the best possible consequences” (Anscombe 1958, p. 9). More

⁵⁰ See e.g. (Glover 1975; Jackson 1987; 1988; Kagan 2011; Parfit 1984; Pinkert 2015; Regan 1980).

precisely, consequentialists hold that possible states of affairs can be ranked, from best to worst, according to their level of an impersonal value ‘goodness’. What I will call ‘pure act-consequentialism’ says that an agent ought to perform an action if and only if it produces the best state of affairs out of those she is in a position to produce. If the agent performs an action which produces any other state of affairs in the ranking, then she fails to do what she ought to do.⁵¹

In this chapter I have committed myself to the ‘sane’ view that consequences matter in ethical discourse. However I have said nothing which commits me to consequentialism. I have not assumed that there is such a thing as impersonal goodness. Furthermore, even given that there is such a thing, I have not assumed that the impersonal goodness of states of affairs determines what ought to be done. To illustrate, consider a Hi-Lo game with two players, x and y , who can each perform an action of either one of two act-types: ‘High’ and ‘Low’. The players face four possible combinations of actions; both do ‘High’, both do ‘Low’ or one does ‘High’ while the other does ‘Low’. Suppose the best outcome will be produced if x and y both do ‘High’. A version of pure act-consequentialism which posits plural reasons will say that x and y ought to do ‘High’ severally. I am not committed to this conclusion. That some action(s) will maximise impersonal goodness might merely be an intelligible reason for some agent(s) to act. As such, it may be a matter of personal taste whether or not we choose to maximise goodness.⁵²

4.5. Conclusion

In this chapter I have argued that progress can be made in resolving certain puzzle cases if we assume that agents in such cases plurally have a reason to act. In coordination problems and cases like Jackson’s driving example, agents plurally have a reason to perform the required combinations of actions. Given this, it is possible that they are open to an instance of moral appraisal on the

⁵¹ See e.g. (Müller 1977; Scheffler 1988; Thomson 2001; Wiggins 2006, chap. 8) for helpful discussions of consequentialism.

⁵² For similar discussions see (Anscombe 1958, p. 12; Munoz-Dardé 2005, pp. 202-203; Pink 2007, p. 425).

basis of their actions. That is, they may be collectively responsible for their actions on (PR). In no difference cases, agents plurally have a reason not to perform some actions which produce an objectionable outcome. As such, they may also be collectively responsible for their actions on (PR). The arguments of this chapter have aimed to establish that there are cases (possible and actual) in which it is plausible to posit the existence of plural reasons. This helps to strengthen the argument of Chapter 3 that collective responsibility as plural responsibility is consistent with responsibility as self-governance. Given that there are plural reasons, it is plausible that there is a plural mode of practical reasoning. As such, there may be a form of self-governance which corresponds to this mode of practical thought.

Conclusion

This thesis has addressed what I have called the problem of collective responsibility; can we understand collective responsibility by employing only the conceptual resources from a theory of singular responsibility? In answering this question I have adopted the following method. I assumed a theory of singular responsibility which I called responsibility as self-governance. I then distinguished three forms of collective responsibility and addressed whether each was consistent with the account of responsibility as self-governance. My enquiry has revealed the assumptions necessary for a deflationary (that is, affirmative) response to the problem of collective responsibility. These assumptions, I suggest, are not wildly implausible. They include the existence of plural actions, the existence of collective agents and the existence of plural reasons. Various authors have affirmed each of these assumptions. However one need not affirm all of them in order to endorse a deflationary response to the problem of collective responsibility. For instance, if one denies that there are such things as collective agents, then one might think that corporate responsibility is inconsistent with responsibility as self-governance. However one may still affirm that both several responsibility and plural responsibility are consistent with this theory.

Various issues have arisen during this investigation which I have not been able to properly address. Most of these issues arise from Chapters 3 and 4. I will mention some of them here with the hope that they will be addressed more thoroughly in the future.

I ended Chapter 3 with a discussion of what I called Sellars' Problem; how does one reason from the assumption that some agents plurally have a reason to act? I did not provide a solution to this problem, though I have hinted at a promising direction (see note 36). Developing such a solution is one obvious area for further work.

A second area for further work is on the topic of holding responsible. I mentioned in Section 1.1.2 that the notions of being responsible and holding responsible are closely connected. Given this, it might seem that the argument of this thesis requires a discussion of how one might hold some agents collectively responsible. This seems clearest in the case of several responsibility; one may resent, blame or punish some agents severally. It may also be unproblematic in the case of corporate responsibility.⁵³ But what about plural responsibility? How does one resent, blame or punish some agents whilst it does not follow that one is resenting, blaming or punishing any one of them?⁵⁴

I will mention one more question which arises from Chapters 3 and 4; why should an agent care about plural responsibility? That is, why should I care whether I am one of a plurality of agents who have failed to act in accordance with plural reasons? After all, it won't follow that *I* will have failed exercise *my* capacity for self-governance as I should.

An initial response to this question is that it may be little different from the traditional question of moral scepticism; why be moral? To illustrate, consider some runners who are competing in the 4 × 100 metre relay. One might ask why any one of the runners should care whether her team wins. After all, each can truly say 'if we win, *I* will not have won'. This thought seems off the mark. We just do care about joint, as well as singular, achievements. The question 'why should I care whether my team wins?' thereby seems little different to the question 'why should I care whether I win?' Perhaps a similar thought holds in the case of responsibility. Maybe people just do care about being one of some agents who have observed, or violated, moral principles. If so then the question 'why should I care whether I am one of some agents who are plurally responsible?' might be little different from the sceptic's question 'why should I care whether I am responsible?'

⁵³ To punish a complex entity may be to dissolve it (Shockley 2007, p. 452). For a discussion of corporations and resentment see (Silver 2005).

⁵⁴ Writers such as Arendt (1987, p. 43) and Young (2011, pp. 76-78) seem to think that one cannot blame some agents plurally. For a discussion to the contrary see (Smith 2009, pp. 54-57).

Bibliography

Anscombe, G. E. M. (1958). Modern Moral Philosophy. *Philosophy*. 33 (124), pp. 1-19.

Anscombe, G. E. M. (1967). Who is Wronged? *The Oxford Review*. 5, pp. 16–17.

Arendt, H. (1987). Collective Responsibility. In: Bernhauer, J. W. (ed.) *Amor Mundi: Explorations in the Faith and Thought of Hannah Arendt*. Boston: Martinus Nijhoff. pp. 43-50.

Boolos, G. (1984). To Be is to be a Value of a Variable (or to be Some Values of Some Variables). *The Journal of Philosophy*. 81 (8), pp. 430-449.

Boolos, G. (1985). Nominalist Platonism. *The Journal of Philosophy*. 94 (3), pp. 327-344.

Bratman, M. (1992). Shared Cooperative Activity. *The Philosophical Review*. 101 (2), pp. 327-341.

Collins, S. (2013). Collectives' Duties and Collectivization Duties. *Australasian Journal of Philosophy*. 91 (2), pp. 231-248.

Corlett, J. A. (2001). Collective Moral Responsibility. *Journal of Social Philosophy*. 32 (4), pp. 573–584.

Cullity, G. (2000). Pooled Beneficence. In: Almeida, M. (ed.) *Imperceptible Harms and Benefits*. Dordrecht: Kluwer Academic Publishers. pp. 1-23.

Deitz, A. (2016). What We Together Ought to Do. *Ethics*. 126 (4), pp. 955-982.

- Descombes, V. (2014). *The Institutions of Meaning: A Defence of Anthropological Holism*. Schwartz, S. A. (trans.) Cambridge: Harvard University Press. Originally published as *Les Institutions du Sens* (1996).
- Feinberg, J. (1968). Collective Responsibility. *The Journal of Philosophy*. 65 (21), pp. 674-688.
- Feldman, F. (1980). The Principle of Moral Harmony. *The Journal of Philosophy*, 77 (3), pp. 166-179.
- Frankfurt, H. (1971). Freedom of the Will and the Concept of a Person. *The Journal of Philosophy*. 68 (1), pp. 5-20. Reprinted in: Watson, G. (ed.) (2003). *Free Will*. 2nd ed. Oxford: Oxford University Press. pp. 322-336.
- French, P. (1979). The Corporation as a Moral Person. *American Philosophical Quarterly*. 16, (3), pp. 207-215.
- French, P. (1982) Crowds and Corporations. *American Philosophical Quarterly*. 19 (3), pp. 271-277.
- Gilbert, M. (2013). *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press.
- Glover, J. (1975). It Makes No Difference Whether or Not I Do It. *Proceedings of the Aristotelian Society, Supplementary Volumes*. 49, pp. 171-209.
- Gold, N. and Sugden, R. (2007). Collective Intentions and Team Agency. *The Journal of Philosophy*. 104 (3), pp. 109-137.
- Held, V. (1970). Can a Random Collection of Individuals be Morally Responsible? *Journal of Philosophy*. 67 (14), pp. 471-481.

- Higginbotham, J. (1998). On Higher-Order Logic and Natural Language. *Proceedings of the British Academy*. 95, pp. 1-27.
- Hossack, K. (2000). Plurals and Complexes. *The British Journal for the Philosophy of Science*. 51 (3), pp. 411-443.
- Jackson, F. (1987). Group Morality. In: Pettit, P., Sylvan, R. and Norman J. (eds.) *Metaphysics and Morality: Essays in Honour of J. J. C. Smart*. Oxford: Basil Blackwell. pp. 91-110.
- Jackson, F. (1988). Understanding the Logic of Obligation. *Proceedings of the Aristotelian Society, Supplementary Volumes*. 62, pp. 255-283.
- Jaspers, K. (2000). *The Question of German Guilt*. Ashton, E. B. (trans.) New York: Fordham University Press. Originally published as *Die Schuldfrage* (1947).
- Julius, A. J. (2013a). The Possibility of Exchange. *Politics, Philosophy & Economics*. 12 (4), pp. 361-374.
- Julius, A. J. (2013b). *Reconstruction*. Unpublished draft. UCLA.
- Kagan, S. (2011). Do I Make a Difference? *Philosophy and Public Affairs*. 39 (2), pp. 105–141.
- Kutz, C. (2000). *Complicity: Ethics and Law for a Collective Age*. Cambridge: Cambridge University Press.
- Lawford-Smith, H. (2015). What We? *Journal of Social Ontology*. 1 (2), pp. 225-249.
- Lewis, H. D. (1948). Collective Responsibility. *Philosophy*. 23 (84), pp. 3-18.
- Linnebo, Ø. (2003). Plural Quantification Exposed. *Noûs*. 37 (1), pp. 71-92.

- MacBride, F. (2005). The Particular-Universal Distinction: A Dogma of Metaphysics? *Mind*. 114 (455), pp. 565-614.
- Massey, G. (1976). Tom, Dick, and Harry, and All the King's Men. *American Philosophical Quarterly*. 13 (2), pp. 89-107.
- May, L. (1992). *Sharing Responsibility*. Chicago: The University of Chicago Press.
- McKay, T. J. (2006). *Plural Predication*. Oxford: Clarendon.
- Müller, A. W. (1977). Radical Subjectivity: Morality versus Utilitarianism. *Ratio*. 19 (2), pp. 115-132.
- Munoz-Dardé, V. (2005). The Distribution of Numbers and the Comprehensiveness of Reasons. *Proceedings of the Aristotelian Society*. 105 (1), pp. 191-217.
- Narveson, J. (2002). Collective Responsibility. *The Journal of Ethics*. 6 (2), pp. 179-198.
- Nefsky, J. (2011). Consequentialism and the Problem of Collective Harm: A Reply to Kagan. *Philosophy and Public Affairs*. 39 (4), pp. 364–395.
- Nefsky, J. (2015). Fairness, Participation and the Real Problem of Collective Harm. In: Timmons, M. (ed.) *Oxford Studies in Normative Ethics, Volume 5*. Oxford: Oxford University Press. pp. 246-271
- Nefsky, J. (n.d.). How You Can Help Without Making a Difference. Unpublished draft.
- Oliver, A. and Smiley, T. (2001). Strategies for a Logic of Plurals. *The Philosophical Quarterly*. 51 (204), pp. 289-306.

Oliver, A. and Smiley, T. (2013). *Plural Logic*. Oxford: Oxford University Press.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Pettit, P. and List, C. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.

Pettit, P. and Schweikard, D. (2006). Joint Actions and Group Agents. *Philosophy of Social Sciences*. 36 (1), pp. 18–39.

Pink, T. (2004). Moral Obligation. *Royal Institute of Philosophy Supplement*. 54, pp. 159-85.

Pink, T. (2007). Normativity and Reason. *Journal of Moral Philosophy*. 4 (3), pp. 406-431.

Pinkert, F. (2014). What We Together Can (Be Required to) Do. *Midwest Studies in Philosophy*. 38 (1), pp. 187-202.

Pinkert, F. (2015). What If I Cannot Make a Difference (and Know It). *Ethics*. 125 (4), pp. 971-998.

Rayo, A. (2002). Word and Objects. *Noûs*. 36 (3), pp. 436-464.

Regan, D. (1980). *Utilitarianism and Cooperation*. Oxford: Oxford University Press.

Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge: Harvard University Press.

Scanlon, T. M. (2003). The Value of Choice. In: Watson, G. (ed.) *Free Will*. 2nd ed. Oxford: Oxford University Press. pp. 352-371. Originally published (1988).

- Scheffler, S. (ed.) (1988). *Consequentialism and its Critics*. Oxford: Oxford University Press.
- Schlick, M. (1939). *Problems of Ethics*. Rynin, D. (trans.) New York: Prentice-Hall. Originally published as *Fragen der Ethik* (1930).
- Searle, J. (1990). Collective Intentions and Actions. In: Cohen, P., Morgan, J. and Pollack, M. E. (eds.) *Intentions in Communication*. Cambridge: MIT Press. pp. 401-415.
- Sellars, W. (1980). On Reasoning About Values. *American Philosophical Quarterly*. 17 (2), pp. 81-101.
- Shiffrin, S. (n.d.). The Moral Neglect of Negligence. Unpublished draft.
- Shockley, K. (2007). Programming Collective Control. *Journal of Social Philosophy*. 38 (3), pp. 442-445.
- Silver, D. (2005). A Strawsonian Defense of Corporate Moral Responsibility. *American Philosophical Quarterly*. 42 (4), pp. 279-293.
- Sinnott-Armstrong, W. (2005). It's Not My Fault: Global Warming and Individual Moral Obligations. In: Sinnott-Armstrong, W. and Howarth, R. (eds.) *Perspectives on Climate Change: Science, Economics, Politics, Ethics*. Burlington: Emerald Group Publishing. pp. 293-315.
- Smart, J. J. C. (1961). Free Will, Praise and Blame. *Mind*. 70, (279), pp. 291-306. Reprinted in: Watson, G. (ed.) (2003). *Free Will*. 2nd ed. Oxford: Oxford University Press. pp. 58-71.
- Smith T. (2005). What is the Hallé? *Philosophical Papers*. 34 (1), pp. 75-109.
- Smith T. (2009). Non-Distributive Blameworthiness. *Proceedings of the Aristotelian Society*. 109 (1), pp. 31-60.

Smith, T. (2011). Playing One's Part. *Review of Philosophy and Psychology*. 2 (2), pp. 213-244.

Spiekermann, K. (2014). Small Impacts and Imperceptible Effects: Causing Harm with Others. *Midwest Studies in Philosophy*. 38 (1), pp. 75–90.

Stone, M. (2001). The Significance of Doing and Suffering. In: Postema, G. (ed.) *Cambridge Studies in Philosophy and Law*. Cambridge: Cambridge University Press. pp. 131-182.

Strawson, G. (1994). The Impossibility of Moral Responsibility. *Philosophical Studies*. 75 (1), pp. 5-24. Reprinted in: Watson, G. (ed.) (2003). *Free Will*. 2nd ed. Oxford: Oxford University Press. pp. 212-228.

Strawson, P. F. (1962). Freedom and Resentment. *Proceedings of the British Academy*. 48, pp. 1-25. Reprinted in: Watson, G. (ed.) (2003). *Free Will*. 2nd ed. Oxford: Oxford University Press. pp. 72-93.

Sugden, R. (1993). Thinking as a Team: Towards an Explanation of Non-selfish Behaviour. *Social Philosophy & Policy*. 10 (1), pp. 69-89.

Sverdlik, S. (1987). Collective Responsibility. *Philosophical Studies*. 51 (1), pp. 61-76.

Thomson, J. J. (2001). *Goodness and Advice*. Princeton: Princeton University Press.

Thompson, D. (2004). The Problem of Many Hands. In: Thompson, D. (2004). *Restoring Responsibility: Ethics in Government, Business, and Healthcare*. Cambridge: Cambridge University Press. pp. 11-32.

Tuomela, R. and Miller, K. (1988). We-Intentions. *Philosophical Studies*. 53 (3), pp. 367-389.

- Uzquiano, G. (2004). Plurals and Simples. *The Monist*. 87 (3), pp. 429-451.
- Uzquiano, G. (2011). Plural Quantification and Modality. *Proceedings of the Aristotelian Society*. 111 (2pt2), pp. 219-250.
- van Inwagen, P. (1990). *Material Beings*. Ithaca: Cornell University Press.
- Velleman, J. D. (1997). How to Share an Intention. *Philosophy and Phenomenological Research*. 57 (1), pp. 29-50.
- Wallace, R. J. (1994). *Responsibility and the Moral Sentiments*. Cambridge: Harvard University Press.
- Watson, G. (1975). Free Agency. *The Journal of Philosophy*. 72 (8), pp. 205-220. Reprinted in: Watson, G. (ed.) (2003). *Free Will*. 2nd ed. Oxford: Oxford University Press. pp. 335-351.
- Weinrib, E. (2013). *The Idea of Private Law*. 2nd ed. Oxford: Oxford University Press.
- Wiggins, D. (2006). *Ethics: Twelve Lectures on the Philosophy of Morality*. Cambridge MA: Harvard University Press.
- Wringe, B. (2010). Global Obligations and the Agency Objection. *Ratio*. 23 (2), pp. 217-231.
- Young, I. M. (2004). Responsibility and Global Labor Justice. *Journal of Political Philosophy*. 12 (4), pp. 365-388.
- Young, I. M. (2011). *Responsibility for Justice*. Oxford: Oxford University Press.