

# **The Scope of Dependence-based Duties**

Stephanie Clare Collins

August 2013

A thesis submitted for the degree of Doctor of Philosophy  
of The Australian National University



## Declaration

The work in this thesis is original and completely my own. No part of this thesis has been submitted for any degree or is currently being submitted for any other degree. All help received in preparing this thesis, and all sources used, have been acknowledged.

Excluding footnotes, this thesis is 94,907 words in length (computer-generated word count).

Signed\_\_\_\_\_ Date\_\_\_\_\_

## Acknowledgements

I arrived in Canberra three-and-a-half years ago with little background in analytic philosophy. I quickly felt out of place and out of my depth. I can't give enough thanks to the people who helped me to find my feet.

To begin, without the encouragement of Paul Daniels back in Auckland, I might not have embarked on a Ph.D. in analytic philosophy. His support in the first several months was invaluable. Then there are those who offered me their camaraderie and intellectual input as fellow aspiring (or recently crowned) ANU Ph.D.s, in particular Alma Barner, Rachael Brown, Ryan Cox, John Cusbert, Adrian Currie, Jonathan Farrell, Luara Ferracioli, Ben Fraser, Jon Herington, Ole Koksvik, Jo Lau, Holly Lawford-Smith, Leon Leontyev, Dan Marshall, Hatha McDivitt, Kelvin McQueen, Dan Schuurman, Rosa Terlazzo, and Clas Weber.

Along the way, these chapters have benefited from the comments of those who attended various seminars and conferences: the Moral, Social, and Political Theory seminar series at ANU (2011); the PhilSoc seminar series at ANU (2010); the Kioloa Graduate Philosophy Conference at ANU (2010; 2012); the Essex Autonomy Project seminar series at the University of Essex (2012); the Political Theory seminar series at the University of Essex (2012); the annual “Brave New World” graduate political theory conference at the University of Manchester (2012); the conference “Collective Intentionality VIII” at the University of Manchester (2012); the graduate philosophy seminar series at the University of Melbourne (2010); the Australasian Postgraduate Philosophy Conference at Monash University (2011); the graduate work-in-progress philosophy seminar series at Monash University (2010); the Australasian Association of Philosophy conference at the University of Otago (2011); the British Society for Ethical Theory conference at the University of Stirling (2012); and the Australasian Association of Philosophy New Zealand Division conference at the University of Waikato (2010). For comments on those occasions and for other conversations about my work, there are too many people to thank individually. Some are named in the previous and next paragraphs. Others are John Broome, Cameron Brown, Tim Dare, Lina Eriksson, Fabian Freyenhagen, Jeff Howard, Larry May, Justin Oakley, Avia Pasternak, Philip Pettit, Felix Pinkert, David Rodin, Ben Saunders, Henry Shue, Holly Smith, Rob Sparrow, Zofia Stemplowska, Christine Swanton, and David Wiens.

For feedback on written chapter drafts and the papers that were connected to them, I thank Elizabeth Ashford, Brett Calcott, Ryan Cox, David Estlund, Jonathan Farrell, Luara Ferracioli, Miranda Fricker, Karolina Milewicz, Rekha Nath, Jonathan Quong,

Frederico Zuolo, and particularly my panel members, Christian Barry and Seth Lazar, as well as Holly Lawford-Smith, with whom I have shared numerous enlightening conversions and email threads about collective duties. My written work was also read and dissected by members of the Nuffield Political Theory Workshop at Oxford University (2012) and the Moral, Social and Political Theory Graduate Workshop at the ANU (2012).

Most especially, this thesis has benefitted from the feedback of my panel chairs, Nic Southwood and Bob Goodin. Nic's ability to tear apart the most carefully constructed argumentative edifice has left me in despair more times than I care to remember, but he has had me up and building again with (I hope) stronger results. Bob has been the guiding star from the day I first proposed the thesis topic. His knack for seeing the big picture, pulling up arguments from disparate and deep wells of knowledge, telling me how to cut out the crap, and doing it all within 48 hours on a 30,000 word draft has never failed to astound me.

Last but not least, I am overwhelmingly grateful to Jonathan Farrell, for adeptly handling my many thesis-induced meltdowns and for his general brilliance at philosophy and life.

## Abstract

If you are dependent on someone in the right way, then that person has a moral duty. We can call this a “dependence-based duty.” This thesis explores the scope of these duties—the range of normative assertions they produce—in both individual and group ethics.

Part I builds a theoretical framework for these duties. I start by developing what I call the “Dependency Principle,” which gives us an ecumenical account of when these duties arise for agents acting independently of others. Roughly, they arise when three conditions are met: (1) the agent is *sufficiently likely* to fulfil an important interest if she takes her most efficacious measures for doing so; (2) the cost of those measures is below some *threshold*; and (3) the cost of those measures is *no higher* than that of any other agent’s most efficacious measures for fulfilling that interest.

This is relatively straightforward for individuals—but some of the most pressing dependence-based duties seem to be held by groups. I give an account of group agents’ capacities to bear duties, before arguing that the Dependency Principle cannot capture the full range of moral duties that intuitively arise out of our dependence on groups. Specifically, it cannot capture cases where no group agent exists, but where a number of individuals could mutually respond to one another to fulfil an important interest, or could come to constitute a group agent that would (if it existed) bear a duty under the Dependency Principle. For these cases, we need a “Coordination Principle” in order to capture all the dependence-based duties that common sense tells us exist. Thus, it is my argument that there are two distinct types of dependence-based duties: one derived from the “Dependency Principle” and the other from a separate “Coordination Principle.”

I then draw out some important implications of the Dependency and Coordination Principles, within interpersonal and international ethics. Part II considers a case from interpersonal ethics. I argue that my two dependence-based principles together provide a unified explanation of the moral theory of care ethics—the result being that care ethics is more systematised than its proponents often claim, and that the dependence-based principles gives us a wider range of claims about interpersonal ethics than first meets the eye. Part III turns to a case from international ethics. I argue that the Dependency and Coordination Principles together give a unified explanation of the international political doctrine of the Responsibility to Protect—the result being that we can more neatly assign the duties of this doctrine than has hitherto been recognised. By seeing that the Dependency and Coordination Principles provide a grounding for the claims of care ethicists and Responsibility to Protect proponents, we can see dependence-based duties as

potentially underpinning many more normative assertions in interpersonal and international ethics than is commonly imagined.

# Table of Contents

Chapter One: Introduction .....	15
PART I – THEORY .....	
Chapter Two: Dependency Duties.....	28
Chapter Three: Collectives' Duties and Coordination Duties .....	61
PART II – INTERPERSONAL ETHICS: THE CASE OF CARE ETHICS.....103	
Chapter Four: The Claims of Care Ethics .....	104
Chapter Five: Unifying and Explaining Care Ethics .....	138
PART III – INTERNATIONAL ETHICS: THE CASE OF RESPONSIBILITY TO PROTECT .....	
Chapter Six: The Agency of States and Intergovernmental Organisations.....170	169
Chapter Seven: Responsibility to Protect: Claims and Potential Explanations .....	187
Chapter Eight: Unifying and Explaining Responsibility to Protect .....	215
Chapter Nine: Conclusion .....	240
Works Referenced .....	245

# Analytical Table of Contents

Chapter One: Introduction .....	15
1.1 The Intuition and the Project .....	15
1.2 Situating the Investigation .....	19
1.2.1 Theories, Principles, Issues .....	19
1.2.2 The Relevant Type of Dependence .....	21
1.2.3 Dependence-based Duties and Capacity-based Duties .....	23
1.2.4 Cognate Projects .....	23
1.3 Thesis Structure .....	25
PART I – THEORY .....	27
Chapter Two: Dependency Duties .....	28
2.1 Introduction .....	28
2.2 Existing Principles .....	29
2.3 “Important Interest” .....	34
2.4 “Sufficiently Capable” .....	38
2.5 “Not Too Costly” .....	42
2.5.1 Sources of Value .....	42
2.5.2 Aggregative and Iterative Value .....	46
2.5.3 A Ranking Method .....	48
2.5.4 Packaging Interests .....	49
2.5.5 The Principle So Far .....	51
2.6 “Least Costly” .....	52
2.7 The Dependency Principle .....	55
2.8 Conclusion .....	59

Chapter Three: Collectives' Duties and Coordination Duties .....	61
3.1 Introduction.....	61
3.2 Duty-bearing Groups .....	63
3.2.1 Agency .....	63
3.2.2 Random Aggregates.....	65
3.2.3 Responsive Individuals.....	65
3.2.4 Collectives .....	70
3.3 Collectives' Duties and Coordination Duties .....	75
3.3.1 Collectives' Capacities and Dependency Duties.....	75
3.3.2 Coordination Duties: The Problem in General .....	80
3.3.3 The Problem for Duties to be Responsive.....	81
3.3.4 The Problem for Duties to Transform Collectives.....	84
3.3.5 The Problem for Duties to Form Collectives.....	86
3.3.6 Solving the Problem: Ruling out some Possibilities.....	87
3.3.7 The Coordination Principle .....	89
3.3.8 Failure, Defection, and Inducing Compliance.....	92
3.3.9 Too Many Coordination Duties? .....	93
3.4 Distributing Collectives' Duties.....	94
3.4.1 Distributing Collectives' Duties to Members.....	94
3.4.2 Collectives' Duties are Made up of Individual Duties.....	97
3.4.3 Collectives' Diachronic Duties and Capacities .....	98
3.4.4 Failure of Collectives and Responsive Individuals.....	99
3.5 Conclusion .....	101

PART II – INTERPERSONAL ETHICS: THE CASE OF CARE ETHICS.....	103
--	-----

Chapter Four: The Claims of Care Ethics .....	104
4.1 Introduction.....	104

4.2 Care Ethics: The General Picture .....	105
4.3 Scepticism about Principles .....	108
4.3.1 Deliberation and Rightness .....	108
4.3.2 Deliberation .....	110
4.3.3 Rightness .....	112
4.4 The Special Moral Value of Personal Relationships .....	114
4.4.1 Initial Characterisation .....	114
4.4.2 The Relevant Relationships .....	116
4.5 Caring Attitudes .....	120
4.5.1 Caring About .....	120
4.5.2 Caring About: The Target .....	122
4.5.3 The Moral Value of Caring About .....	123
4.5.4 Duties to Care About .....	126
4.5.5 The Analysis .....	128
4.6 Caring Actions .....	129
4.6.1 Caring For: The Target .....	129
4.6.2 The Moral Value of Caring For .....	133
4.6.3 The Analysis .....	135
4.7 Conclusion .....	136
 Chapter Five: Unifying and Explaining Care Ethics .....	138
5.1 Introduction .....	138
5.2 Initial Evidence for a Dependence-like Ground .....	139
5.3 Claim 1: Sympathetic and Contextually-orientated Deliberation .....	144
5.3.1 The Principles Call for Some Sympathy and Contextuality .....	144
5.3.2 The Principles Call for Enough Sympathy and Contextuality .....	146
5.3.3 The Principles Explain Claim 1 .....	148
5.4 Claim 2: Relationship Importance .....	150
5.4.1 The Principles Generate Some Relationship Duties .....	150

5.4.2 The Principles Generate Enough Relationship Duties .....	152
5.4.3 Extension: Duties to Form Personal Relationships.....	155
5.4.4 The Principles Explain Claim 2.....	157
5.5 Claim 3: Caring About .....	159
5.5.1 The Principles Call for Caring About .....	159
5.5.2 The Principles Call for Enough Caring About .....	160
5.5.3 The Principles Explain Claim 3.....	162
5.6 Claim 4: Caring For .....	164
5.6.1 The Principles Call for Caring Actions .....	164
5.6.2 The Principles Call for Enough Caring Actions.....	165
5.6.3 The Principles Explain Claim 4.....	166
5.7 Conclusion .....	167

### PART III – INTERNATIONAL ETHICS: THE CASE OF RESPONSIBILITY TO PROTECT .....

Chapter Six: The Agency of States and Intergovernmental Organisations.....	170
6.1 Introduction.....	170
6.2 States.....	170
6.2.1 States as Agents .....	170
6.2.2 Collective Members of Collectives .....	173
6.2.3 Individual Members: False Negatives and False Positives? .....	173
6.2.4 Autochria.....	176
6.2.5 Democracia .....	177
6.2.6 The Need for Non-state Collectives .....	178
6.3 Inter-governmental Organizations .....	180
6.3.1 The General Picture.....	180
6.3.2 IGOs as Real Collectives .....	183

6.4 Conclusion.....	186
Chapter Seven: Responsibility to Protect: Claims and Potential Explanations.....187	
7.1 Introduction .....	187
7.2 The Canonical Formulations of the Three Claims.....	192
7.2.1 ICISS 2001 .....	192
7.2.2 UNGA 2005 .....	199
7.2.3 Ban 2009 .....	201
7.3 Possible Justifications and Precisifications .....	203
7.3.1 Self-interest.....	203
7.3.2 Human Rights .....	205
7.3.3 Voluntary Assumption.....	208
7.3.4 Contribution.....	209
7.3.5 Association .....	211
7.3.6 Proximity.....	212
7.4 Conclusion.....	213
Chapter Eight: Unifying and Explaining Responsibility to Protect.....215	
8.1 Introduction .....	215
8.2 Clarifying the Harms and Actions .....	216
8.3 The Domestic Duty .....	219
8.3.1 The Dependency Principle as Necessary and Sufficient.....	219
8.3.2 Complications with Domestic Duties.....	223
8.4 International Duties: Single-agent Cases.....	224
8.4.1 Initial Evidence for a Dependence-like Ground.....	224
8.4.2 The Dependency Principle as Sufficient .....	226
8.4.3 The ICISS's Criteria for Intervention.....	227
8.5 International Duties: Multi-agent Cases.....	232
8.5.1 Calls for Coordination in the Literature.....	232

8.5.2 The UN as Sole Duty-bearer?.....	234
8.5.3 The Coordination Principle as Sufficient .....	235
8.6 Conclusion .....	238
Chapter Nine: Conclusion .....	240
Works Referenced .....	245

# Chapter One: Introduction

## 1.1 The Intuition and the Project

This thesis is about dependence. It clarifies and builds upon the following basic thought: if you are dependent on someone in the right way, then that person has a moral duty. We can call this a “dependence-based duty.” The aim is to specify what “dependent on someone in the right way” means, and then see what normative demands this basic thought generates in interpersonal and international ethics.

Let us start with a variation on a well-worn example. You are sunbathing at the beach when you notice a young child in trouble in calm water. Your companions are mediocre swimmers, while you are a very able swimmer. Or perhaps you are each equally able to swim, but the rescue would be much more costly for each of them. Or you are each equally able to rescue the child at similar cost, but any one of them would probably injure the child in the process, while you would not. In these cases, there is a dependence relation between you and the child. Roughly, that is to say: the child has an *important interest* that is unfulfilled and you are *sufficiently able* to fulfil that interest at *not too high a cost* to yourself and the child—and not only this, but your best attempt to rescue the child would realise *less cost* to yourself and the child than any other agent’s best attempt to do so. We can gloss this kind of dependence by saying that you are “best-placed” to fulfil the child’s important interest.

To see that this relationship is plausibly duty-generating, stipulate the following: you realise the trouble the child is in and that you are best-placed to help. Yet suppose you sit and watch the child drown. Have you have defaulted on a duty? It certainly seems so.<sup>1</sup> And the weight of this duty seems to be grounded not only in your being minimally *able* to rescue the child (though that is certainly necessary). Rather, your duty carries heavy weight—your defaulting seems particularly blameworthy—because you are *best-placed* to rescue him. The fact that you are best-placed gives you, out of everyone on the beach, weighty reason to help.

Similar duties seem to be held by groups, as well as individuals. Suppose the child is far out in the rip tide, and cannot be rescued by any one person alone. The beach is

---

<sup>1</sup> Of course, not all philosophers would agree. In discussing duties to assist, Jan Narveson (2003), for example, claims that you have done nothing *wrong* by defaulting on those duties. But even he agrees you have done something *very bad*.

unpatrolled by surf lifesavers but, luckily enough, the off-duty members of a surf lifesaving team happen to be amongst the crowd on the beach that day. These people can work together to save the child. No team member could do this alone, but they have a well-established procedure for responding to such situations. By implementing the procedure, they will almost certainly save the child, at very low cost to themselves and to the child. Perhaps it is possible that some other random collection of beachgoers would save the child if they worked together, but they are not as likely to succeed as the team. Or perhaps the random collection is as likely to succeed as the team, but the random collection's trying would have higher costs for themselves or for the child than the team's trying would (say, they would save the child but be exhausted, or they would injure him in the process). In this case, the *team* seems to be best-placed, and so the *team* seems to have a particularly weighty duty (at least in the first instance: how the duties of collectives distribute to their members is a further question). And the team seems to have this duty even though they are off work: the duty is grounded not in the expectations of their job, but in the fact that they are best-placed.

Formal groups like this are relatively clear-cut, but there are far more difficult cases involving groups. Imagine that there are numerous strangers at the beach, and the child cannot be rescued by any one of them acting alone, but he can be rescued through the combined efforts of the five best swimmers present. Again, perhaps other random collections of beachgoers would have some chance of saving the child. But imagine that these other collections would have a greater risk of injuring the child, or would incur greater cost to themselves by trying. The first collection—the five best swimmers—do not know each other and have no established procedure for saving the child. Nonetheless, they seem (in some sense) together to be best-placed to rescue him. If the situation were transparent to them, it would be wrong for them not to try. Yet there is no team at which we can direct our thanks, if the child is saved, or our anger, if he isn't: there are no teams on the beach, only collections of strangers. Accounting for the precise distribution of duties in such “ad hoc” group cases is difficult. But this thesis takes seriously the conviction that there are duties in such cases.

These cases are stylised and abstract. Yet the relevant kind of dependence is not confined to philosophers’ thought experiments. It permeates our lives. When we are infants, we depend on others to feed, clothe, and carry us; when we are sick, we depend on others to nurture, mend, and calm us; and when we are old, we depend on others to comfort, nurse, and console us. Often there are many people who *could* do these things for us. But we tend to turn first to those whose resources, or relationship with us, renders them best-placed to help us.

And these kinds of dependence hold even for those of us lucky enough not to live in a constant state of hunger, persecution, or deprivation. A great many human beings are dependent in much more long-term and much more systemic ways, for example as a result of blind political hatred, corrupt governments, or insufficient resources. They depend on their families and local communities, but also on their states, and on the complex workings of intergovernmental organisations, for their basic interests in subsistence and security to be met. If we take seriously the intuitive appeal of the stylised beachside rescue duties—as held by both individuals and groups—then this will have ubiquitous real-world implications, both interpersonally and internationally.

It is therefore surprising that, despite these duties' intuitive appeal and real-world ubiquity, their upshots remain incompletely explored, in both theory and practice, and at both the individual and group levels. Regarding the *theory* of *individual* dependence-based duties, we need to know: over which interests do these duties range? Can there be dependence-based duties to provide others with anything whatsoever, or just the most important things? How should we understand the ability, or capacity, that is required in order to have one of these duties? What kinds of costs constrain these duties? How exactly must the duty-bearer's capacity and costs compare to that of other agents—what precisely does it mean to be “best-placed,” or to have someone dependent on *you* in particular?

Concerning the *theory* of *group* dependence-based duties, we might ask about the possibility and nature of group agency, group capacities, and group duties: in what sense might a group of individuals be best-placed to provide assistance, if none of them alone can exercise the capacity to assist? Can we talk about “the group” exercising the capacity (and having a duty to do so), and, if so, what does that mean? How do groups' duties distribute to individuals? And if only certain kinds of groups (say, those with moral agency) are eligible to bear duties, then what duties might there be in cases where no such group exists but a number of individuals could create a best-placed duty-bearing group, or could together meet some vitally important interest?

About the *practice* of *individual* dependence-based duties, we might wonder: should we really be consciously entertaining dependence-based duties as we go about everyday life, or are they somehow self-effacing? How many of the duties recognised by common sense can be understood as instances of them? And can these duties call for attitudes and emotions, as well as actions?

The *practice* of *collectives*' dependence-based duties is equally complex. Can important real-world groups, such as states and intergovernmental organisations, bear dependence-based duties? If so, what does this imply for their members? To what extent are these

duties borne out in actual political practice? How might they be fulfilled in international politics?

This thesis answers these questions in a way that helps to shed light on various moral phenomena. In Part I (Theory), I explore the theory of individual and group dependence-based duties. I answer, or give frameworks for answering, the two sets of theoretical questions listed above. I explain the conditions under which there exist dependence-based duties for agents acting alone, whether those agents are individual or collective (I call these “dependency duties”), as well as the conditions for dependence-based duties for agents acting together who do not constitute a collective agent (I call these “coordination duties”). Dependency duties and coordination duties are the two types of dependence-based duty with which the remainder of the thesis is concerned.

In Parts II and III, I argue that an appeal to Part I’s theoretical account enables us to unify, precisify, and explain two highly suggestive but unsystematised normative doctrines. In Part II (Interpersonal Ethics), I consider the moral doctrine of care ethics. In Part III (International Ethics), I consider the political doctrine of Responsibility to Protect (R2P).<sup>2</sup> Within their particular domains of concern, these doctrines give answers to, respectively, the individual practical and collective practical questions that I listed above. I will use my theoretical analysis to make sense of the answers that these two doctrines give to the practical questions with which they are each concerned. In using my theory to generate the practical normative claims of only two doctrines, my aim is to give a narrow, but deep, account of my theory’s implications.

These two doctrines are largely unexplored by analytic moral philosophers. This is regrettable, as each doctrine encapsulates some important insights about the practical implications of dependence-based duties. By showing how these duties can generate the core claims of these two doctrines, my argument serves a dual function. First, it advances care ethics and R2P, by anchoring their normative claims with an intuitive and precise category of duties. Second, it advances more general theorising about dependence-based duties’ scope, by elaborating upon these duties’ implications within care ethics’ and R2P’s respective domains of concern. As will be seen, these two doctrines cover a range of moral and political issues, and give complex solutions to a range of problems. By seeing the normativity of dependence as behind these solutions to these problems, we will see that dependence-based duties can give us much more of personal and political morality than is commonly recognised.

---

<sup>2</sup> As we shall see, care ethicists address institutions and international ethics, as well as individuals and interpersonal ethics. Yet I call care ethics “interpersonal” because the care ethical focus remains firmly on natural individuals, rather than the artificial collective agents that my group theory—and R2P—deal with.

## 1.2 Situating the Investigation

### 1.2.1 Theories, Principles, Issues

In conceptualising the kind of philosophy at play in this thesis, it is helpful to think of different levels of moral theory. Normative ethicists and political philosophers are often concerned with examining foundational theories: in normative ethics, theories like consequentialism, contractualism, and virtue ethics; in political philosophy, theories like liberalism, cosmopolitanism, and communitarianism. These foundational theories of the right or good (ethics), or the just (politics), are theories about what makes things right, good, or just, simpliciter and in all circumstances. These occupy the bottom, foundational layer of normative space. At other times, normative ethicists and political philosophers are concerned with resolving specific issues on which foundational theories bear: in ethics, issues like abortion and vegetarianism; in politics, issues like taxation and immigration. These occupy the surface (or at least, a near-surface) layer of normative space.<sup>3</sup> Dependence-based duties occupy an intermediate layer of normative space, between the theories and the specific issues to which those theories apply. I will call this the layer of *principles*.<sup>4</sup> Principles can be justified in terms of foundational theories and can be used to justify judgments on specific issues.

Principles are a useful layer of enquiry: they are apt to garner consensus and they are explanatory. Philosophers—not to mention everyone else—are more likely to be able to agree on principles than on foundational theories. As we shall see, philosophers from various foundational camps can agree that something like dependence-based duties exist. Moreover, one can engage in theorising at the level of principles even if one is sceptical about the whole project of foundational theorising (as Lichtenberg (2004) does). Yet unlike directives about what our duties are in particular cases (such as arise at the level of issues), principles tell us *why* we should do what we should do. They are more explanatory than judgments about particular issues, even if they lack the full depth of explanation found at the foundations. This mixture of being both consensus-apt and explanatory makes principles ripe for philosophical exploration.

Principles take the form of conditionals. They state that if certain conditions obtain, then certain things ought to be the case. I will be concerned in particular with principles in

---

<sup>3</sup> There are plausibly more than three levels. Above the issue of abortion, for example, we have the issue of *this woman's* abortion.

<sup>4</sup> The label “principle” is stipulative: of course it makes perfect sense to talk about foundational principles, like “the right action is the one that maximises happiness,” or “individual liberty is the only political value.”

which the “ought” is a duty. There are numerous examples of this type of principle, many of which are likely entailed by more than one of the ethical and political theories I mentioned above. There is the contribution principle (if one has contributed to harm, then one has a duty to provide remedy for that harm) (Pogge 2002); the promise principle (if one has induced another to rely on one’s performing some action, then one has a duty to perform that action) (Scanlon 1998, ch. 7); the beneficiary principle (if one has benefited from an injustice, then one has a duty to compensate the victims of that injustice) (Butt 2007); the association principle (if one is in an associative relationship with somebody else, one has a duty to give special consideration to them) (Sandel 1982; Scheffler 1997; Scheffler 2001); and so on. Alongside these sits the dependence-based principle: roughly, if an agent (or set of agents) is best-placed to meet someone’s important interest, then the agent (set) has a duty to do so.<sup>5</sup>

The consequent of a duty-generating principle will contain either an “all-things-considered” duty or a “pro tanto” duty. We can think of pro tanto duties as very weighty—though not overriding—moral reasons. We can think of all-things-considered duties as what we have moral reason to do once all the pro tanto reasons are taken into account. If we frame principles in a pro tanto way, then all pro tanto principles whose antecedents apply in a given case must be weighed up when we are figuring out what is morally required in that case all-things-considered. For example, suppose I have promised my mother that I will phone her tonight, but then I knock a colleague down the stairs and he needs taking to hospital. Whether I should phone my mother or take my colleague to hospital (assuming I cannot do both) will depend on how we should weigh up the promise principle and the contribution principle (as well as, perhaps, a capacity principle, discussed below). This is a question on which the principles themselves are silent, for which we must turn to our foundational theories.

Principles can instead be formulated in an all-things-considered way, though their antecedents will then need to mention other principles. It will be necessary to include in the antecedent “and no other stronger moral reason applies” (or some such). Again, whether this part of the antecedent holds will be a question for the foundational theory. Regardless of whether we go “pro tanto” or “all-things-considered,” the moral considerations in each principle need to be counterbalanced with each other when we are establishing what we have a duty to do all-things-considered. (As we shall see, in formulating my principles I will

---

<sup>5</sup> Rough variants are found in: Goodin 1985; Scanlon 1998; Singer 1972; Singer 2009; Unger 1996. These are discussed in §2.2.

opt for an all-things-considered formulation with a condition that no defeating reason applies. Nothing of substance hangs on this.)

### **1.2.2 The Relevant Type of Dependence**

It will also be helpful at the outset to distinguish the sense of “dependence” I have in mind from related, but distinct, notions. We sometimes use “dependent” as a one-place predicate: someone might be “dependent” without depending *on* anything or anyone. This is to say they are deprived, helpless, or in an otherwise wretched state, without assuming that anyone or anything could respond to that state. This is the sense in which “dependent” contrasts with “independent.” A baby on an empty desert island might be dependent in this sense.

I am not concerned with this kind of dependence. I am concerned with the possibility of *responding* to dependence, so I focus upon dependence as a relation between a dependent and an entity *on which* they depend.<sup>6</sup> To depend on someone or something is to need them or be vulnerable to them. Moreover, because I am concerned with dependency *duties*, I am concerned with cases where the latter entity is a moral agent. The sense in which the farmer depends on the weather, or the addict depends on her drug, will not concern me. (Though the sense in which the farmer depends on the weather forecaster, and the addict depends on her dealer, have not yet been ruled out.) My concern with duties also means I am concerned with cases where the dependent is an entity with moral status—an entity that deserves consideration in moral theorising and deliberation.

Yet dependence as a two-place relation—“X depends on Y”—is also not my ultimate concern. I am primarily interested in the three-place relation: “X depends on Y *for* Z.” Specifically, I focus on cases where X depends on Y for the fulfilment of one or more of X’s “interests,” in the broadest possible sense of that term. So I will be concerned with dependents that have interests. Entities with interests are “moral persons” as I will use the term. Of course, the fulfilment of interests is not necessarily the only source of moral value—it is just the source of moral value with which this thesis is concerned. (For ease of exposition, I will further limit my discussion to human beings, but this isn’t meant to rule out dependency duties to fulfil the interests of other beings.)

---

<sup>6</sup> Corresponding to the *Oxford English Dictionary*’s third definition of “dependence”: “The relation of having existence hanging upon, or conditioned by, the existence of something else; the fact of depending *upon* something else.” (Of course, mine isn’t a claim about existence. It will instead be a claim about the fulfilment of important interests.)

There are related notions that we should disentangle from dependence. First, to say that X is dependent on Y for Z is not necessarily to evoke counterfactual dependence, in the metaphysicians' sense. That is, X might depend on Y for Z even if it is possible that X will come to enjoy Z by some method that has nothing to do with Y. For example, we might say that a five-year-old depends on her parents for food, even though it is possible—perhaps even probable—that she will find food in some other way if her parents fail. Perhaps a kind neighbour will step in, or perhaps the state will. So the child's having food does not *necessitate* that her parents do anything. Nonetheless, there is something about her parents that picks them out in particular, and that makes it correct to say that she *depends* on them for food. (Exactly what this “something” is will occupy me in Chapters Two and Three.)

A second related notion is *reliance*. When we rely on someone, we expect, believe, trust, or at least intend to act as if they will behave in a certain way. As I am concerned with it, dependence does not necessarily include such psychological states. In fact, dependence does not necessarily include *any* psychological states: you might depend on someone without any awareness of it, just as the patient in a coma might depend on the doctor without any awareness of it.<sup>7</sup>

Third, to depend on someone is not necessarily to be subjugated, subservient, or subordinated to them. Of course, with dependence often comes power. Because A depends on B *for* something, B's capacity to withhold that something gives B power with respect to A. But dependence is not necessarily inferiority. The power does not necessarily permeate the dependence relationship in general. As we will see in Chapter Five, some of the most important dependence relations are entirely symmetrical, such as those between some spouses.

Fourth, dependence also not entail that your life plans, emotions, or projects are tied up in a constitutive and ongoing way with the person on whom you depend. Dependence can be a very short-term, unlikely, and incidental state of being—as the happenstance-like beach rescue example demonstrates.

---

<sup>7</sup> In this way, I am concerned with a more narrow breed of dependence or vulnerability than that grounded in *reliance* or *expectation*, which is part of the notion of “vulnerability” that Robert E. Goodin uses to analyse the responsibilities promisers have to promisees (1985, 44), friends have to friends (1985, 97), and at least some responsibilities beneficiaries have to benefactors (1985, 103). These moves will not work with the kind of dependence on which I am focused (though in Chapter Five I will show that my narrower conception of dependence does have upshots for friends).

### **1.2.3 Dependence-based Duties and Capacity-based Duties**

It is also worth distinguishing dependence-based duties from another, more well-established, class of duties: capacity-based duties. Roughly, the capacity principle states that if one *can* fulfil someone's important interest (perhaps at not-disproportionate cost to oneself or others), then one has a duty to do so (D. Miller 2001, 460–461). Dependence-based principles can be understood as a type of capacity-based principle, insofar as their driving force derives from the ability to fulfil an important interest and the fact that the exercise of that ability will not cost too much. But dependence-based duties have an additional feature, namely, that the relevant agent's (or agents') most efficacious measures for fulfilling the important interest will realise *no less cost* than any other agent's most efficacious measures for doing so. That is, dependence-base duties are about being *best-placed*. (In Chapter Two, I will add more flesh to these bones, including other subsidiary necessary conditions.)

Why write a thesis about dependence-based duties in particular rather than capacity-based duties in general? There are at least three reasons. First, the duties of the best-placed are arguably more weighty than other capacity-based duties, such as those that accrue to any and all agents that merely meet some threshold of capacity. This makes dependence-based duties an important breed of capacity-based duty. Second and relatedly, I suspect there is more consensus about the existence of dependence-based duties than there is about capacity-based duties as a whole. This is because dependence-based duties are less ubiquitous, since they are borne by only one agent (or, as we shall see, a relatively small subset of all capable agents) in relation to a given interest. Third, in using just this subset of capacity-based duties to explain care ethics and R2P, I set myself a more ambitious task in Parts II and III of the thesis. Explaining these doctrines using the more general class of duties would arguably be too easy, since there would be scope to draw on so many different breeds of capacity-based duties to explain different parts of the doctrines.

### **1.2.4 Cognate Projects**

Finally in situating my investigation, it is worth noting that other philosophers have engaged in projects that could be described as exploring the scope of capacity-based duties, just as this thesis explores the scope of dependence-based duties. By seeing how the current project differs from these others' projects, the aims and scope of this project will be more clearly demarcated.

First, Peter Singer (1972; 2009) and Peter Unger (1996) take our intuitions about duties grounded in capacities, and use them to demonstrate that morality demands much more of us than we ordinarily think, particularly with regard to alleviating global poverty. While I have no objections to Singer's or Unger's projects, mine is different, in at least three ways. First, as just noted, they deal with duties based on one's merely being capable of fulfilling another's interests, rather than on one's being best-placed. Their evocation of these duties is thus possible in a broader range of contexts than the evocation of dependence-based duties would be.

Second, they aim to draw *counter-intuitive yet true* conclusions about the practical upshots of capacity-based duties, while I aim to show that dependence-based duties entail and justify conclusions that are, by and large, intuitive. I aim to show that many intuitive claims about individual and group morality—claims that are already endorsed by care ethicists and R2P proponents—can be *precisified, unified, and justified* using duties that derive from being best-placed. The idea that these claims can be cleaned up and explained using dependence-based duties is novel; the claims themselves are not.

Third, neither Singer nor Unger give a systematic account of collective duties, or duties of individuals to coordinate in fulfilling others' important interests. When Singer comes to impasses at which duties might be construed individually or collectively, he explicitly focuses on what any given individual should do on their own (e.g., 2009, 39, 53–56). And Unger (1996) does not even mention the possibility of collective duties.<sup>8</sup> An account of collectives' duties is an important contribution of the current project.

My project also has close affinities with that of Robert E. Goodin (1985). Goodin argues that the standard repertoire of “special responsibilities”—the responsibilities we have to our “families, friends, clients, and compatriots”—are best-explained by the fact that they are “particularly vulnerable to our actions and choices.” Others are also vulnerable to us, for example “foreigners, future generations, animals, and natural environments” (1985, 11, 186). Therefore, there are vulnerability-based duties to the latter as well as the former. This project differs from mine in at least three ways.

First, like Singer, Goodin's primary aim is to increase the number of duties that we recognise (1985, 9), rather than to demonstrate that important existing doctrines are explicable by a sufficiently detailed account of dependence-based duties. Thus the counter-

---

<sup>8</sup> At one point Unger notes that, in the context of international aid, individual beneficence works in a causally “amorphous” way, since “on one end of a causal chain, there are many donors contributing together and, on the other, there are all the people saved by the large effort they together support.” But, he says, “since there's nothing morally objectionable about proceeding to aid greatly needy folks amorphously, no moral weight attaches to the precise character of the causal relations between the well-off and those whom, whether collectively or not, they might help save” (1996, 48–49).

intuitive duties that Goodin explains are, by and large, very different from the intuitive ones I aim to explain (indeed, the doctrines I aim to unify under dependence-based duties either did not exist, or barely existed, in 1985).

This difference points to a second difference between my account and Goodin's: our notions of vulnerability/dependence. In using vulnerability to explain some already-recognised duties, Goodin sees *reliance* and *expectation* as important kinds of vulnerability. The dependence-based duties at issue in this thesis, however, are not of that kind, so the duties that it explains—and thus the subject matter it deals with—will be correlatively different from Goodin's (see fn. 7).

Third, Goodin's move from already-recognised duties to counter-intuitive duties crucially requires imposing duties on groups. Here, Goodin assigns responsibilities to organised and disorganised groups indiscriminately (1985, 68, 136–9, 151). A large part of my project is concerned with being very precise about the nature of group dependence-based duties, and how they operate in agent versus non-agent groups. If my analysis is right, Goodin's move from the recognised duties to the unrecognised ones requires much more fine-grained analysis and argument than he provides.

### 1.3 Thesis Structure

The structure of the thesis will be as follows. In Chapter Two I will lay out a detailed conditional called the “Dependency Principle.” The development of this principle constitutes the *individual theoretical* component of the thesis. The principle states the conditions under which there are dependency duties (i.e., dependence-based duties for moral agents who are acting independently of others), and explains many of the dependence-based duties that seem intuitively to exist. According to it, dependency duties exist where one agent is best-placed to fulfil an important interest. An agent is best-placed when, roughly, she is the agent whose most efficacious measure for fulfilling an important interest meets three conditions: it is *sufficiently likely* to fulfil the important interest; it has *low enough* expected cost; and it will realise *no more* expected cost than the most efficacious measure of any other agent.

In Chapter Three, I tackle the crucial *group theoretical* issues. I argue that only groups with a certain kind of structure—call them “collectives”—can bear duties, and I provide a reductively individualistic account of collectives’ duties. Collectives, where they exist, can bear dependency duties under the Dependency Principle. However, I show that the Dependency Principle cannot capture the full range of moral duties that seem to arise out of our dependence on groups: sometimes no one individual or collective agent is best-

placed, but the fulfilment of some important interest nonetheless hinges on the intertwined actions of several individuals. In such cases, the Dependency Principle needs supplementing with a “Coordination Principle” in order to capture all the dependence-based duties that intuitively exist. The latter principle constitutes an important contribution to more general theorising about individuals’ duties in group contexts, as well as to dependence-based duties in particular.

Chapters Two and Three constitute Part I: Theory. The two principles developed in Part I—the Dependency and Coordination Principles—together give rise to the full range of dependence-based duties with which this thesis is concerned. The remaining chapters argue that they can give us a compelling unifying account of two separate doctrines.

In Part II (Interpersonal Ethics; Chapters Four and Five), I address the *individual practical* issues as they play out within the ethics of care. The Dependency and Coordination Principles are used to precisify, unify, and explain the claims about interpersonal ethics that are made by care ethicists. These include the claims that principles should play a marginal role in moral deliberation, that personal relationships have an underappreciated moral status, and that ethical behaviour includes an expansive range of attitudes and actions. Chapter Four develops a statement of the four key claims of care ethics, explaining how this version of care ethics arises naturally out of critical reflection on the literature. Chapter Five explains how the Dependency and Coordination Principles give us those four claims. Chapter Five starts by giving some textual evidence for a dependence-like explanation of care ethics, before going through each of the care ethical claims in turn, explaining how my two principles provide enlightening explanations of them.

Part III (International Ethics; Chapters Six, Seven, and Eight) explores some *group practical* upshots of my theoretical analyses. Chapter Six segues into Part III, by explaining how states and intergovernmental organisations are eligible for duties under the Dependency and Coordination Principles. Chapter Seven identifies the three core claims of the emerging international political doctrine of R2P—first, that states have duties to protect their populations from certain harms, while, second, other states and international agents have duties to assist states in protecting their populations and, third, other states and international agents have duties to step in if any state fails. I demonstrate the current lack of clarity about R2P’s moral basis and practical implications, before considering and expressing doubts about possible unifying explanations of the doctrine. Chapter Eight then argues that the Dependency and Coordination Principles can precisify, unify, and explain the three claims of R2P.

Chapter Nine concludes the thesis.

## **PART I – THEORY**

## Chapter Two: Dependency Duties

### 2.1 Introduction

Dependency duties are dependence-based duties for agents acting independently of others. The basic thought behind these duties is this: if you are dependent on an agent in the right way, then they have a moral duty. This basic thought has a conditional structure. The current chapter investigates this conditional. The aim is to arrive at an interpretation of “dependent on an agent in the right way” that makes such dependence duty-generating—that is, an interpretation of the antecedent that makes the conditional true. The interpretation aims to capture as many intuitions as possible, through a process of reflective equilibrium between our intuitions about scenarios and the features those scenarios have in common.

To that end, I start in §2.2 by surveying what other philosophers have done to precisify the basic thought, or thoughts like it. This will help us to see the issues my principle will have to address. To anticipate, the general structure of the relation that gives rise to dependency duties is this. A is dependent (in the relevant sense) on B just in case four rough conditions are met: A has an *important interest* that is unfulfilled; B is *sufficiently capable* of fulfilling that interest; B’s most efficacious measure for fulfilling the interest will be *not too costly*; and B’s most efficacious measure for fulfilling the interest will be the *least costly* of anyone’s most efficacious measure for doing so.

§2.3–2.6 engage in in-depth discussion of each of these four components: “important interest,” “sufficiently capable,” “not too costly,” and “least costly.” I address each of these notions in a way that captures the intuitive pull of the basic thought. I aim to explain them in a way that remains as neutral as possible between the different possible justifications or broader theories that underlie the basic thought of which these four components are each a part. Thus in §2.7, we arrive at the Dependency Principle: a conditional that provides an ecumenical, yet detailed, account of the conditions under which there are dependency duties.<sup>9</sup>

---

<sup>9</sup> Much of what I say in this chapter applies to dependence-based duties for agents acting together, as well as for agents acting independently of others. But some of what follows is peculiar to the latter cases. For that reason, I will talk in the remainder of this chapter only of “dependency duties,” rather than of “dependence-based duties,” where the former arise in “acting alone” cases and the latter encompasses “acting alone” and “acting together” cases. In Chapter Three, I will have recourse to refer back to some of the notions developed in this chapter for addressing dependence-based duties for agents acting together.

## 2.2 Existing Principles

Other philosophers have offered principles that are intended to capture capacity-based duties. As we saw in Chapter One, dependency duties can be thought of as a type of capacity-based duty. So, one might think that my principle for dependency duties can piggyback on other theorists' principles for capacity-based duties. We might just take other theorists' principles for capacity-based duties, and tweak them so they are about the duties of the best-placed agent who is acting independently of others. Unfortunately, the existing principles are insufficiently precise to be of much use. By quickly reviewing the gaps in these principles, we will be in a position to see which details mine has to capture.

First, Singer has given three different capacity principles<sup>10</sup>: “if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it” (1972, 231); “if it is in our power to prevent something very bad from happening, without thereby sacrificing anything morally significant, we ought, morally, to do it” (1972, 231); and “if it is in your power to prevent something bad from happening, without sacrificing anything nearly as important, it is wrong not to do so” (2009, 15).

Each of these principles is insufficiently specified. They don't give us tools for specifying how much power is power enough, how bad a bad is bad enough, or which sacrifices we need to measure the importance of. As it happens, Singer's substantive views mean that his principles generate extremely counterintuitive duties. According to Singer (1972), we have duties to help others, insofar as we can, until helping any more would make us worse off than those we are trying to help. Assuming these are all-things-considered duties (which Singer seems to imply they are), it should be possible to formulate the principle in a way that either gives a role for other principles to defeat these demands, or that better explains just why the duties are so very demanding.<sup>11</sup> (I will opt for the former.) Additionally, Singer's principle does not say how we should prioritise our resources if there are multiple bads we could prevent. And it does not consider the possible

---

<sup>10</sup> Scott James (2007) uses the fact that these are “mere capacity” principles to criticise Singer's argument. James notes a difference between the “drowning child” example (which Singer uses to motivate his principles) and Singer's principles: in the drowning child case the agent is *uniquely* capable of saving the child. Singer uses this example to generate principles that deal with cases where when one is *merely* capable. James views this move as illegitimate. My principle will not suffer from James' objection: I will use cases where one is not *uniquely* capable, but rather *best-placed* of all agents present; and my principle will likewise apply to cases where the agent is *best-placed* (not merely capable or (necessarily) uniquely capable).

<sup>11</sup> Though Singer does hint at a place for other principles when he admits that “‘nearly as important’ is a vague term. ... I don't know what *you* might think is as important, or nearly as important, as saving a life. ... I'll trust you to be honest with yourself about that” (2009, 17).

trade-offs between acting alone and acting in concert with others. Almost certainly, Singer's aim was to be imprecise so as to garner consensus (as he suggests in Singer 2009, 17). But my upcoming analysis will allow us to retain the consensus while getting more precise about how "bad," "power," "sacrifice," and so on relate to one another.

Peter Unger (1996) skirts around the project of stating precise principles, though he makes one "fallible formulation of a fair bit" of his moral view:

Insofar as they need her help to have a decent chance for decent lives, a person must do a great deal for those few people, like her highly dependent children, to whom she has the most serious sort of special moral obligation. Insofar as it's compatible with that, which is often very considerably indeed, and sometimes even when it's not so compatible, she must do a lot for other innocent folks in need, so that they may have a decent chance for decent lives. (1996, 12)

This principle sounds like it has two incompatible consequents: doing "a great deal" for a "few people" *and* doing "a lot for other innocent folks in need." (Fred Feldman (1999) interprets the principle like this and criticises it on that basis.)

But upon closer inspection, the principle states that we must do a great deal for these special "few people" (such as our children) *insofar as* they need our help to have decent chances for decent lives. This consideration—having a decent chance of a decent life—is exactly what drives our obligations to "other innocent folks in need," according to Unger. So the two obligations look to be the same, with the principle really being: "do a great deal (i.e., a lot) to help people get decent chances for decent lives (perhaps sometimes prioritising those who depend on you and you alone, such as your children—but perhaps not)." Seen in this way, the principle looks like a consequentialist one—with the good to be maximised being a decent chance of a decent life. Yet Unger elsewhere denies his account is consequentialist (Unger 1999, 213), leaving this detail unclear.

Whether or not Unger's principle is consequentialist, it leaves unaddressed the possibility of other considerations weighing against the injunction to give people decent chances of decent lives. (Though Unger does—very briefly—endorse obligations to financially support children and spouses, which he seems to think exist over and above this principle (1996, 149–150).) Other crucial matters are also left out of the principle. How much is a "great deal"? Should this "great deal" take account of previous sacrifices? In other work (1999, 203–6), Unger suggests it should not—but this is intuitively highly problematic, as Hooker (1999, 180–181) notes. Must we only act to give people decent chances of decent lives, or do we have correlative weaker (or stronger) reasons to give

them poor chances of decent lives, decent chances of amazing lives, amazing chances of decent lives, and so on?

Most importantly, since my interest is in developing a principle for the *best-placed* agent, how might we use Unger's principle to *compare* various agents' capacities, sacrifices, and previous actions? Unger's principle contains no tools for doing this. (Of course, this was not Unger's aim. So this is not a criticism of Unger, *per se*. It is rather a criticism of the suggestion that we might build on his principle in formulating a principle for the duties of best-placed agents.)

A third formulation of the capacity principle is given by Goodin's "first principle of individual responsibility": "[I]f  $A$ 's interests are vulnerable to  $B$ 's actions and choices,  $B$  has a special responsibility to protect  $A$ 's interests; the strength of this responsibility depends strictly upon the degree to which  $B$  can affect  $A$ 's interests" (1985, 118). That is, the *strength* of  $B$ 's duty is proportionate to the *amount*  $B$  can affect  $A$ 's interests. Goodin allows that " $B$ 's responsibility for protecting  $A$ 's interests is ... always susceptible to being overridden by  $B$ 's other moral responsibilities" (1985, 118). This leaves it open that the existence and strength of  $B$ 's duty will be affected not only by the amount  $B$  can do for  $A$ , but also by how much  $B$  can do for others.

This principle does not pick out any agent as the duty-bearer: it says only "vulnerable," not "particularly" or "especially" vulnerable. Though Goodin uses the "particularly vulnerable" formulation elsewhere (1985, 11, 33–4), it is not immediately clear what would constitute particular vulnerability, or being best-placed, under an extension of the above principle. It's thus not clear how helpful Goodin's principle could be for achieving our aim in this chapter. His principle suggests that we are particularly vulnerable to those who can affect our interests to the greatest degree, but there are at least four problems with this.

First, what if the agent who can affect my interests to the greatest degree would incur greater costs than other agents in doing so? Would she have the duty of particular vulnerability, or would it rather fall to an agent who can affect my interests less, but at substantially lower cost to herself? We might think "not taking on disproportionate costs" is a responsibility of the agent's (perhaps a responsibility to herself), that can, in accordance with Goodin's view, "override" the vulnerability-based duty. But it seems that whether an agent is "best-placed" should depend not just on her other responsibilities (as Goodin allows), but more centrally on others' abilities to affect the dependent's interests and their costs of doing so.

Second, even in cases where there is just one potential duty-bearer, it seems that *what* or *how much* is demandable of  $B$ —and not just how *stringent* that demand is—should be

proportionate to the degree to which B can affect A's interests. Scope and stringency come apart: the scope (what or how much is demandable) might be very low (say, \$10 rather than \$100), even though the stringency of that demand is very strong (in that it would be very wrong not to meet the demand). If scope and stringency are separate outputs of the principle, we should be able to say how they interact.

Third, Goodin's principle—even as a principle of mere capacity—seems to over-generate duties. As the principle is stated, I have a responsibility for everyone whose interests are causally downstream of my actions and choices. While these responsibilities' strengths vary according to my potential effects on those interests, this still generates more responsibilities than is common sensically plausible. (It's also unclear what the sense is in which these numerous duties can be construed as “special” responsibilities, as Goodin's principle says they are.) If we want our principle to be ecumenical with respect to different foundational theories, we should give scope for the differing limits they would set on when these duties arise. For example, perhaps these duties should only arise when I can affect their interests to a *sufficient degree*, or affect *particularly valuable* interests, or have a good enough *chance* of affecting their interests if I try.

Fourth, Goodin's account does not say enough about how the interests of multiple vulnerable persons might interact with one another in generating my duty. He allows the interests of one person to override the interests of another, but what if I could fulfil more interests of more persons if I “packaged” those interests together? Goodin's account certainly does not forbid this kind of packaging, but it does not positively account for it. It is preferable to have an account that explicitly addresses these issues, rather than one that is silent on them.

A fourth and final formulation of the capacity-based principle comes from T.M. Scanlon. According to Scanlon, “[t]he cases in which it would most clearly be wrong not to give aid—and most clearly unreasonable to reject a principle requiring that aid be given—are cases in which those in need of aid are in *dire straits*: their lives are immediately threatened, for example, if they are starving, or in great pain, or living in conditions of bare subsistence” (1998, 224, emphasis added). Scanlon's foundational theory states that an act is wrong if a principle permitting it is reasonably rejectable, so he endorses the principle that “[i]f you are presented with a situation in which you can prevent something very bad from happening, or alleviate someone's dire plight, by making only a slight (or even moderate) sacrifice, then it would be wrong not to do so” (1998, 224.). This principle is “one that could not reasonably be rejected, at least not if the threshold of sacrifice is understood to take account of previous contributions (so that the principle does not demand unlimited sacrifice if it is divided into small enough increments)” (1998, 224).

Scanlon also recognises that it would be unreasonable for anyone to reject a principle on which we are required to help people when they are not in *dire* need, but where we could help them pursue their life's projects with *no significant* sacrifice to ourselves. In general, Scanlon says, the more the agent can do for the person in need, the higher is the cost that can be demanded of the agent (1998, 227–8).<sup>12</sup>

Scanlon's account contains some useful details. It says something about how bad the bad has to be (it might be dire or not, depending on the sacrifice the agent must make to remedy it) and it mentions which sacrifices of the agent we must take into account (previous ones as well as current ones). I will develop an account that agrees with these points, while adding more. For example, Scanlon hasn't told us how to describe (i.e., pick out) the dire (or not-so-dire) plight that we are remedying, or what it is for an agent to be "able" to prevent a bad or alleviate some plight. And he doesn't tell us how to deal with situations where many agents are each capable of rectifying the plight. The latter situations matter particularly for duties that accrue to the *best-placed* agent, as these duties need to pick out just one agent as the duty-bearer. (But they also matter for duties of mere capacity, if the plight will be alleviated only if *exactly* one capable agent acts to alleviate it, and we need to pick who that one should be.)

These brief comments are not intended to refute Singer's, Unger's, Goodin's, or Scanlon's principles. In fact, they each make important moves in specifying the slippery basic thought behind capacity-based duties. This slippery basic thought is: if you *can* fulfil someone's important interest (at not-disproportionate cost), then you have a duty to do so. The slippery basic thought behind dependency duties is slightly different. The dependency version of the basic thought states that if you are *best-placed* to fulfil someone's important interest, then you have a duty to do so. Unfortunately, these theorists' specifications of the capacity-based thought do not point directly to a specification of the dependency-based thought. We need a more thoroughgoing and detailed account of the notions evoked in the basic thought.

However, to retain these duties' appeal in a way that is consistent with a wide variety of foundational theories, what we do *not* need is a substantive conception of how to judge the relative importance of interests, or of the extent to which personal cost rules out moral duties. That would be to side with one foundational theory at the expense of another. We rather need a framework—a common scaffolding—for understanding the way these notions relate to one another, in order to demarcate the subject matter of

---

<sup>12</sup> Scanlon distinguishes this from the claim that *the worse off the person is*, the higher the cost that can be demanded of the agent. The moral requirement becomes more demanding if the required action would do more for the person. This might not hold in all cases where the person is very badly off.

dependency duties as a type of duty common to a range of foundational theories. Theorists who adhere to different foundational theories will then be able to pin their respective theories of interests, costs, and so on onto that framework. The remainder of this chapter builds such a framework.

As I said in introducing this chapter, the general structure of the relation that gives rise to dependency duties is this. A is dependent (in the relevant sense) on B just in case four rough conditions are met: A has an *important interest* that is unfulfilled; B is *sufficiently capable* of fulfilling that interest; B's most efficacious measure for fulfilling the interest will be *not too costly*; and B's most efficacious measure for fulfilling the interest will be the *least costly* of anyone's most efficacious measure for doing so. This is the key idea, and it will do for some purposes. But not for others. I will now carefully examine each of these four ideas in turn. As I go through these four ideas, the four conditions will be refined and important supplementary conditions will be uncovered.

## 2.3 “Important Interest”

First, we need to know how conceptualise the interests that give rise to dependency duties. I will call these the “important interests.” If we allow ourselves to take the concept of an interest as a primitive, then the question becomes one of how to distinguish the “important” interests from the “unimportant” ones.

It is useful to start by distinguishing instrumental interests from final interests. Dependency duties often require us to fulfil instrumental interests. But dependency duties are always re-describable in terms of final interests, the fulfilment of which happens to require that we fulfil the instrumental interest. This is because the importance of instrumental interests derives entirely from the importance of the final interests with respect to which they are instrumental, taking account of the availability, expected efficacy, and costliness of other instruments.

To appreciate the distinction, consider the following example. If we live in an area where malaria is prevalent, we plausibly have an interest in having mosquito nets. But this is only true because (and insofar as) we have an interest in good health; the former is a mere means to the latter, and any dependency duty to fulfil the former is describable as a duty to fulfil the latter in circumstances under which a mosquito net is an instrument to that fulfilment. Moreover, if there are available instruments to avoiding malaria other than mosquito nets (say, a tablet), then one's interest in a mosquito net might become less important, even though one's interest in good health remains just as important. (Though it might still be vitally important that one gets the net *or* the tablet—or even both, if they

both only reduce the likelihood of malaria rather than ensuring against it, and the interest in good health is extremely important.)

The question for a theory of dependency duties then becomes: which final interests are important enough to generate dependency duties? I propose we imagine a *continuum of final interests*, from trivial to important. Which final interests are further towards the “important” end of the continuum will depend on one’s substantive theory of welfare, about which this thesis is fairly agnostic.<sup>13</sup>

My locution “important interests” is intended to capture the idea that interests are only ever more or less important. Thus “important interest” should be understood to mean “*relatively or more* important interest.” The continuum view allows us to say that the more important the interest, the more morally valuable it is that you have it fulfilled. All else being equal, then, the more important the interest, the more likely there is to be a dependency duty to fulfil it. However, as we shall see below, there might be a dependency duty to fulfil a relatively non-important interest, if an agent is extremely capable of doing so at extremely low cost to herself and the dependent.

Some theories of interests envisage a clear cut-off point on the continuum of interests, beyond which one has stronger, or a new kind of, claim to have interests fulfilled. The interests beyond this cut-off point might be referred to as “vital interests,” “needs,” “basic rights,” “basic entitlements,” or similar. One common view is that these are constituents of any conception of the good life. Using this view, we can include such things as subsistence, security, and liberty at the “important” end.<sup>14</sup>

It is consistent with my account that there are “lexically prior” interests, which must be fulfilled before we worry about other interests, or that there are “particularly valuable” interests, whose fulfilment is *much* more valuable than the fulfilment of any other interests. If so, then the value of fulfilling these interests will trump (for lexically prior interests), or be much higher than (for particularly valuable interests), that of fulfilling lesser interests—at least, all else being equal. If such interests exist, then dependency duties to fulfil these interests are liable to arise much more often than dependency duties to fulfil other interests. But my account neither assumes nor denies the existence of this prioritised category of interests.

---

<sup>13</sup> I cannot be completely agnostic between these, of course. For example, if I allow “most important interest” to mean “most preferred state of affairs”, then my account will generate duties to buy someone an expensive sports car, if having an expensive sports car is their preferred state of affairs, I am best-placed to fulfil this interest, and the costs of fulfilling this interest are low enough. I am agnostic between the *live candidates* for being our best theory of welfare. I assume that the view I just described is not a live candidate.

<sup>14</sup> These are what Henry Shue (1996) identifies as “basic rights.” Others have more or less extensive lists, building on (more or less) the same idea that there are certain final interests that are common to all humans (e.g. Nussbaum 2006, 78; D. Miller 2007, ch. 7).

There are a few more clarifications regarding important interests. First, later on in the thesis, I will assume that certain interests are sufficiently important to generate dependency duties. These include interests in loving relationships and the distinctive kind of care they typically involve (the focus of Part II), and the interest in protection from mass atrocities such as genocide, war crimes, ethnic cleansing, and crimes against humanity (the focus of Part III). In assuming that these interests are important enough to generate dependency duties, I may seem to be assuming that the importance of an interest for a given person should not be understood subjectively, but rather objectively—that is, not “from the point of view of that person’s tastes and interests,” but “independent of that person’s tastes and interests” (Scanlon 1975, 656, 658). It is true that I will talk as if these interests are important for everyone, and in general I will talk as if there is an objective fact about which interests are important. But I do not mean to imply that the importance of all interests is defined objectively, or that there are no subjectively defined important interests.

Second, one might think we have an interest not just in having our interests fulfilled, but in having them fulfilled securely or robustly. If this is true of some interest, we should include it in the description of the interest when we are determining the value of that interest’s fulfilment. We can, for example, consider Alma’s interest in good health given the way things are now, and her interest in having her good health “socially guaranteed against standard threats” (in Shue’s (1996) locution). There might be a duty to fulfil *each* of these interests, perhaps where each duty accrues to a different agent—the first duty might accrue to her doctor, and the latter to her government, for example. The duties over these two interests accrue to different agents. So, for the purpose of describing the interests there are dependency duties over, we should include such provisos within the definition of the interest only on an interest-by-interest basis. This allows us to capture Alma’s doctor’s duty.

Third, when we are considering the continuum of interests—from important to non-important—we must not conflate “important” with “urgent.” It might be extremely important that some interest is fulfilled, but the relevant action need not be taken now. Like “securely,” “with an action taken now” should be added to only some interests. Imagine there is someone whose life I can save if and only if I act today, and there is someone else whose life I can save if and only if I act any time within the coming year. The first case is more urgent than the second, but my reason to act *sometime in the next year* in the second case is just as strong as my reason to act *today* in the first case. If we precisify the interests—indexing one to “fulfilled with an action taken now” and indexing the other to “fulfilled with an action taken up to one year from now”—then the associated duties might

accrue to different agents. One interest is more *urgent* (will need tending to earlier), but neither would be more *important*.

Fourth, some interests might be indexed to the agent who fulfils them. For example, I might have an important interest not just in being loved, but in being loved *by my parents* (Liao 2006; Keller 2006). This is important for dependency duties: if my interest is in being loved by my parents, then only my parents will be eligible for a dependency duty to fulfil that interest. But other agents might have duties to (take measures that might) make my parents capable of loving me. These other people's duties would be duties to fulfil my interest in "having parents that are capable of loving me." Because my interest in having parents that are capable of loving me is instrumental to my interest in parental love, third parties' dependency duties to fulfil the instrumental interest will depend upon the value of the final interest (parental love) and the likelihood that the final interest's fulfilment will follow from the instrumental interest's fulfilment.

Fifth, dependency duties always range over *unfulfilled* interests. But the fact that a person has some good now does not mean that all of her interests related to that good are now fulfilled. For example, my need for "food now" might be fulfilled if I have just eaten. But if I will not have enough food tomorrow, then my interest in "food tomorrow" is not fulfilled. Moreover, dependency duties are always duties to take measures that (are sufficiently likely to) *fulfil* interests—not just, say, that are sufficiently likely to contribute to fulfilling them, or to increase the likelihood that they will be fulfilled. So it matters that an important interest might be, say, "having a better chance of nourishment" or "being better nourished" rather than simply "being well-nourished." This allows that others might have dependency duties to *add to* my nourishment, even if they are incapable of fully nourishing me. By doing so, they would *fulfil* my interest in additional nourishment.

In all these ways, interests can be defined with various caveats and indexes. They will accordingly move up and down the spectrum of importance, and may generate duties for different agents.

Sixth and finally, for the purpose of identifying dependency duties, interests should be defined as generally as possible while capturing all that is of value in them. For example, if what matters is just "that I am fed tonight," then we should not split hairs between my interest in "getting fed tonight by A" and my interest in "getting fed tonight by B." If there is no morally important difference between these interests, then our account should not pointlessly say that they each have a duty over these different interests, based on their respectively being best-placed to fulfil those agent-indexed interests (at least, not if we assume full compliance—a point I will address in §2.6). Assuming they are not equally well-

placed to fulfil my interest in *being fed tonight*, only one of them should have a duty to fulfil that interest on the basis of being best-placed to do so.

The upshot is this. A has a dependency duty to take some measure only if:

- (1) B has an important final interest, FI, that is unfulfilled, where the interest is described in a way that does not include morally irrelevant facts (though it might be temporally or otherwise indexed).

This is clearly a necessary, not sufficient, condition. We should turn to consider the other necessary conditions.

## 2.4 “Sufficiently Capable”

I will assume that dependency duties are always duties to take some measure (whether an action or an omission) that has at least some likelihood of fulfilling a (sufficiently) important interest. I will use the term “measure” in a way that assumes agents have full control over the measures they take.

In order for an agent to have a duty to take a measure *in virtue of that measure’s propensity to fulfil an important interest*, the measure must be sufficiently likely to fulfil that interest. There are issues with how to understand “sufficiently likely,” here. For example, it seems strange to assign to an agent a duty to take some measure, in virtue of that measure’s propensity to fulfil an important interest, if the measure has only a 0.0001 likelihood of succeeding. Yet in some cases—if the interest is important enough—an agent can have a duty to take measures that have a tiny likelihood of fulfilling an interest. For example, if the interest at stake is “all persons’ final interests in being alive next week” and A could take the measure “firing a rocket at the giant asteroid that, if not hit by a rocket, will hit Earth and kill everyone this weekend,” A might have a duty to take those measures, even if the measures are unlikely to fulfil all those final interests. There are measures he can take that will *possibly* fulfil the interests, so he can have a duty to take those measures, where the duty is grounded in the importance of the interests.

I suggest, then, that we say the following: in order for an agent to have a duty to take a measure, where the duty is grounded in an important interest, the likelihood that the measure will fulfil the interest must be *proportionate* to the importance of the interest.<sup>15</sup> The

---

<sup>15</sup> This adapts Geoffrey Brennan and Nicholas Southwood’s (2007) analysis of “ $\varphi$  is feasible for A,” which is “A has a sufficiently high likelihood of  $\varphi$ -ing, conditional on trying.” I use “proportionate” to suggest that the sufficiency threshold can be different in different cases.

less important the interest, the more likely it must be that some measure will fulfil it, in order for the measure's likelihood to be proportionate to the interest's importance. Conversely, for an extremely important interest, a measure might have a low likelihood of fulfilling it (if the measure is taken), and yet it might be proportionate.

For simplicity, let's use “A's likelihood” to mean “the likelihood of the measure of A's that is *most efficacious*—i.e., the likelihood of the measure of A's that is most likely (of all A's measures) to fulfil the interest if A takes it.” And let's use “success” to mean “the important interest is fulfilled.” In order for A to have a dependency duty grounded in some interest, then, A's likelihood must be in the range of likelihoods that are *proportionate* to the importance of the interest, where more important interests are proportionate to a range of likelihoods whose lower bound is lower. If A's likelihood of success is proportionate to the interest, then A is *sufficiently capable* of fulfilling the interest to bear a dependency duty to take the relevant measures.<sup>16</sup> Thus “sufficiently important interest” and “sufficient capacity to fulfil the interest” interact.

One might wonder what we should do if A's likelihood of success depends on other variables. For example, we might ask what A's likelihood is if other agents try to stop him, or if the laws of nature suddenly change. In general, the likelihood of all these sorts of conditions holding is included in A's likelihood. Thus if it is very likely that others will try to stop A, or that the laws of nature will suddenly change to render success impossible, then this reduces A's likelihood of success.

However, we should not allow A's likelihood to be affected by the possibility that *other agents* will work with A towards fulfilling the interest. This is because in these circumstances, it would not, properly speaking, be A that succeeds. To see this, consider a six-year-old's capacity to bake a cake. Imagine that if the six-year-old takes measures to bake the cake, it is very likely that a cake will be baked. This is because his father will read the recipe, measure the ingredients, offer advice, turn on the oven, and check that the cake is done. All the child will actually do is mix the ingredients. Here, there is a very high likelihood of the cake being baked if the child takes the relevant measures, but it rings false that the child can bake the cake. This is because the father must intentionally alter his actions in response to the child, and vice versa: they must work together in an ongoing way, communicating and planning together. To describe the cake-baking as a capacity of

---

<sup>16</sup> This seems to present a problem: is the relevant likelihood here objective probability, subjective probability (if so, whose?), evidence-relative probability (if so, whose evidence?), or what? I take it our best theory (whatever that turns out to be) of moral duties under uncertainty can apply here. Presumably, it will refer to the reasonable beliefs that some actual or ideal agent does or should or would have about the agent's likelihood of  $\varphi$ -ing. For some possibilities, see Smith 2010.

the child elides important facts, namely, facts about the intertwined intentional involvement of another agent.

It is a different story if the father is replaced with a robot that the child is able to manipulate. The robot can be treated as a feature of the child's environment. The father can also be treated as a feature of the environment, if he does not work with the child (i.e., act responsively to him, cooperate with him, mould his actions around the child's in an ongoing way) with the goal of baking the cake. In these scenarios, we can take the robot and the father's actions as fixed: as features of the child's environment to which the child might respond. The likelihood of these environmental features "cooperating" can be factored into the child's likelihood of success. But if the child and the father *work together* in order for the cake to be baked—if they intentionally coordinate as they go—then the likelihood of success here does not accrue to the child alone. Such cases are crucially important for dependence-based duties, and I will deal them in detail in Chapter Three. For now, however, I am concerned with the dependence-based duties of agents acting independently of others, that is, where other agents do not *intentionally coordinate* with the agent with the aim of fulfilling the interest in question. This is because I am concerned for the moment with instances where persons depend upon *one* agent. My reasons for refusing to include coordination cases as dependency duties will become clearer in Chapter Three.

If the exclusion of coordination-based likelihoods seems to reduce dependency duties' potential too much, the next point will restore it. This point is that capacities sometimes require multiple steps to be exercised. Sometimes we are insufficiently likely to fulfil some interest now, but we are able to make ourselves sufficiently likely to fulfil the interest later (Goldman 1978; Jensen 2009; Gilabert & Lawford-Smith 2012). To really lack sufficient likelihood of fulfilling an interest, the agent has to be insufficiently likely to fulfil the interest not just with her currently available most efficacious measures, but also with measures that will be available to her if she takes certain current measures. One can have *diachronic* dependency duties to take a whole string of measures. Indeed, all dependency duties are plausibly diachronic to some extent: in order to pull the child from the shallow water, you might first have to take off your shoes, step into the water, and so on.

Relatedly, we can sometimes be diachronically dependent on others to ensure that they *do not* have certain dependency duties at some later time. For reasons related to other ethical and political principles, or to one's theory of interests, it might be valuable that certain of our interests are fulfilled by ourselves, rather than having others fulfil them for us. And others are sometimes capable of *making us such that* we are capable of (perhaps continuing to) fulfil our own interests. My interest is in "being made such that I can fulfil

my own important interests,” and you might have a dependency duty to do something now so that I can fulfil that interest later.

A final clarification regarding sufficient capacity relates not to fulfilling any particular important interest, but rather to making oneself *likely to be* sufficiently capable in the future. Intuitively, it is not morally permissible to avoid incurring dependency duties by intentionally diminishing one’s capacities with the primary goal of avoiding dependency duties that one believes might arise to fulfil interests of existing persons.<sup>17</sup> Imagine that Peter decides not to seek a cure for his fear of water primarily in order to ensure that, when and if the time comes, he will be able to cite that fear as rendering him incapable of rescuing drowning children. We need to assume the fear is curable, but that the cure would take, say, several hours (so that it cannot be cured when he sees the drowning child in enough time for him to save her). We also need to assume that this fear, if it is not cured in enough time, renders Peter incapable of rescuing the child: it is a chronic and debilitating fear.

If Peter avoided curing such a fear, then he would be morally deficient. Yet it is false that there is an identifiable interest that he is (even diachronically) sufficiently capable of fulfilling, if he has not yet encountered any drowning children, or any other important interests whose fulfilment requires comfort with water. We thus need to allow that there is a supplementary duty over and above dependency duties: a duty to refrain from intentionally diminishing one’s measures’ efficacy primarily in order to avoid dependency duties that might arise. This is not itself a dependency duty, but it is closely related. This duty will re-emerge in §2.7.

Let me recap the two conditions for dependency duties that I have discussed. In order for A to have a dependency duty to take some measure M, it must be the case that:

- (1) B has an important final interest, FI, that is unfulfilled, where the interest is described in a way that does not include morally irrelevant facts (though it might be temporally or otherwise indexed); and
- (2) If: A intentionally takes measure M, where M is the most efficacious measure open to A to fulfil the interest in (1); then: even if other agents do not coordinate with A, FI will be fulfilled with a likelihood that is *proportionate* to FI’s importance,

---

<sup>17</sup> I add the caveat “to fulfil interests of existing persons” in order to allow that it might be permissible for one to avoid creating new persons (i.e., having children) primarily in order to avoid having dependency duties. Non-identity problems notwithstanding, I do not wish here to take a stand on the permissibility of such avoidance.

where more important interests are proportionate to a wider range of likelihoods (with a lower bound that is lower).

## 2.5 “Not Too Costly”

### 2.5.1 Sources of Value

It is typical to claim that we are not obliged to take interest-fulfilling measures if doing so would be excessively costly for ourselves. We call such cost to the agent “personal cost”—it is the cost that accrues to the agent in virtue of his taking some measures.<sup>18</sup> We can think of this as “opportunity cost”: it is whatever the agent might do instead of taking those measures. As we’ve just established, the measures we are specifically concerned with are the agent’s most efficacious measures for fulfilling an important interest. Intuitively, the personal cost of these measures should not be too high.

Yet to talk about costs is to neglect appreciation of the *benefits* that might be realised from the taking of an interest-fulfilling measure. It is therefore more illuminating to talk about *personal value* rather than *personal cost*. The personal value of an agent’s most efficacious measures for fulfilling an important interest is the sum of personal benefits minus the sum of personal costs, given that the agent takes these measures rather than doing whatever else he might do instead. (I consider how to compare one agent with another in §2.6.)

However, the agent is not the only one whose costs and benefits are relevant to the dependence relation. The dependence relation holds between an agent and *another moral person*. So we should also factor in the costs and benefits of the agent’s measure for this moral person—that is, for the dependent. This includes the benefit to the dependent of having the important interest fulfilled, and perhaps other effects on the dependent as well. For example, even if the important interest would be fulfilled by the measures, perhaps numerous other interests of the dependent would be set back by them. The total value of the measures, as regards the dependent, should take account of all of these factors. (If the child is saved, but is knocked around and suffers concussion, then the value of the rescue measures, vis-à-vis the dependent, are diminished relative to a rescue without concussion. If the non-concussion rescue would have happened in the absence of the concussion rescue, then this will be an opportunity cost for the child.)

---

<sup>18</sup> The distinction between personal cost and other kinds of cost is disputable. We can say, for example, that personal cost is *built into* morality, such that it is really a moral cost (Goodin 2009); or that the limits of reasonable personal cost are determined by the presuppositions of morality (Cullity 2009; Cullity 2004a ch. 8). For ease of exposition I will separate personal and other kinds of cost.

That is to say, when we are questioning whether a given interest-fulfilling measure is too costly, we should consider: the benefits to the agent, minus the costs to the agent; plus the benefits to the dependent, minus the costs to the dependent. These costs and benefits are to be measured from the baseline of whatever might have happened instead (weighted for the probability of those happenings). The resultant value might be negative or positive.

Importantly, I do not mean for these “benefits” and “costs” to be read in a strictly utilitarian way—all kinds of intrinsic, extrinsic, final, and instrumental values and disvalues to the agent and the dependent might be included. Often, we will only be able to say whether the value of some measure is positive or negative, or greater or lower than the value of some other measure—we will rarely be able to put a precise number on these values. This is in part because it is plausible that “different values are relevant in different ways in different decision contexts” (Anderson 1993, 47), such that there is no *general* function for determining the value of a measure for the agent and dependent, where that function can be applied to each possible measure in any possible scenario.

And even within a given scenario, it might be that various values (each one arising out of different possible measures the agent could take) can only be given an ordinal ranking, not a cardinal one. Sometimes, even an ordinal ranking might not be possible—there may be no fact of the matter about which measures have greatest value. For example, if A’s most efficacious measure for fulfilling an important interest would have the result that A could not write a great novel (that A otherwise would write), while B’s most efficacious measure for fulfilling that important interest would have the result that B could not paint a great painting (that B otherwise would paint), then there might just be no fact about which person’s measure realises highest cost to the relevant agent—because the novel and the painting are incommensurable (Anderson 1993, ch. 3; Raz 1986, ch. 13).<sup>19</sup>

In general, then, the idea of a measure having “positive” value is somewhat metaphorical: what I mean by a measure having “positive” value is that it is better, regarding the agent and the dependent, that the measure is taken than that it is not.

---

<sup>19</sup> This is perhaps a less-than-optimal result, but I assume that if one believes values are incommensurable, then one will accept such results for a whole range of moral duties. In other words, this problem—that if not all values can be placed on one scale, then the moral upshots are unclear—is not particular to dependency duties. Moreover, the compelling examples of dependency duties (such as beach rescue cases) suggest that this incommensurability does not always paralyse our ability to make judgments about which measures have positive value. And as we shall see in §2.6, my analysis of “best-placed” will state that the best-placed agent is that whose most efficacious measures will *not realise less* value than any other agent—it will not require that she *realise more* value than any other agent. On this analysis, all else being equal, both the novelist and the painter would have a dependency duty.

One's judgment about the value (for agent and dependent) of a measure will largely be determined by one's theory of interests and by the other kinds of values one endorses. For example, some foundational theories might allow us to weight costs to the agent higher than costs to the dependent, or vice versa. One way they might do the former is by saying that the inability to exercise personal prerogatives greatly increases the costs to the agent of taking an interest-fulfilling measure. If one holds such a view, then one will more often judge the value of an agent's most efficacious interest-fulfilling measures to be negative.

Even more strongly, one's foundational theory might have a prohibition on certain costs to the agent—his life or his arm, say. One might believe that no duty can possibly demand that the agent make such a sacrifice. If so, then that cost serves a “blocking” function: when losing an arm is placed on the “cost” side of the scales, no benefits may be added on the other side (or they can be added, but it is impossible that they will outweigh the cost). This blocking function will result in a negative value for any measure that involves the agent losing his arm—*regardless* of the benefits to the dependent.<sup>20</sup> (As we shall see, this would rule out the possibility of a dependency duty to take such a measure.) Whether such “blocking” costs exist is not a question on which my account of dependency duties will take a stand.

One question about these values on which I will take a stand arises from the fact that dependency duties are forward-looking imperatives, not backwards-looking assessments of value. In order for us to determine whether there is a *duty* to take some measure, the costs and benefits (for agent and dependent) that might be realised by the measure have to be weighted by their *likelihoods*. If some measure will in fact realise great value, but it is not reasonably expectable—given the evidence that is available—that the measure will realise that value, then we should not say that the agent who fails to take that measure has done something wrong. For this reason, when it comes to determining whether the value realisable by some measure rules out a dependency duty to take that measure, we must consider its expected value, not its actual value.

In sum, the “expected value” of a measure is equal to: the potential benefits for agent and dependent, given that the measure is taken, multiplied by the benefits’ likelihoods, minus the potential opportunity costs for agent and dependent, given that the

---

<sup>20</sup> On personal cost constraints on capacity-based duties, see: Kumar 1999; Ashford 2003; Cullity 2004a Part II; R. Miller 2004; Barry and Øverland 2013; Sonderholm 2012. The issue of the demandingness of morality more generally comes up most prominently in discussions of consequentialism (Sobel 2007; Singer 2009), but it arises for other foundational theories such as contractualism (Ashford 2003). Regardless of foundational theory, it is usually our duties to assist the needy—of which dependence-based duties are one type—that are the target of various “demandingness” objections to morality, and to the idea that there are some costs morality cannot demand of agents (see, e.g., Feltham and Cottingham (eds) 2010; Chappell (ed.) 2009).

measure is taken, multiplied by those costs' likelihoods. As long as this value is positive—i.e., above 0—the costs are not disproportionate to the benefits, so the costs are insufficient to render the action non-obligatory. That is, a necessary condition for a dependency duty to take a measure is that the measure's expected value (regarding agent and dependent) is positive. Precisely which costs and benefits matter, and how they are variously weighted, is a question for foundational theories.

One might think that the idea of “expected value” allows too much to get in the way of the basic intuition behind dependency duties. The basic intuition was about my taking measures to fulfil some *specific important interest*. If we allow all other costs and benefits to myself and the dependent get into the picture, then we have missed the point that this was meant to be about fulfilling that particular important interest. The rule that “the expected value regarding agent and dependent must be positive” seems to be equivalent to a satisficing consequentialist injunction to make sure expected value is positive with regard to those two persons.

But note that a measure's having a positive expected value is merely a *necessary* condition for a dependency duty. It is not sufficient. We do not have dependency duties to take just *any* measures that have positive expected value with regard to ourselves and someone else—we only have *dependency* duties to take such measures if there is a specific important interest that the measures are sufficiently likely to fulfil, and are our most efficacious measures for fulfilling. But if those measures will realise more costs than benefits for the relevant parties—where the benefits include the value of the important interest being fulfilled—then it is natural that this should block the possibility of a duty.

Additionally, the clause that the measures must be “not too costly”—or, in the more precise phrasing, must “realise positive expected value regarding agent and dependent”—does not yet allow us to say that the resultant duty is an *all-things-considered* duty. This is because all other moral and non-moral principles and values have not been factored into this condition. Only costs and benefits for the agent and dependent have been factored in, as only these are relevant for the question of whether there is a *dependence relation*. A dependence relation, as we saw in Chapter One, holds between three entities: an agent, a moral person, and an interest of the moral person. So the only values relevant to whether this relationship holds are values to the agent and the moral person (where value to the moral person includes the value of the interest being fulfilled).

To see that other values are plausibly not part of the dependence relation, consider the following case. Walter has a good chance of saving Jesse from breaking his ankle. As far as Walter and Jesse are concerned, Walter's measures have positive value. However, in exercising those measures, Walter is sure to kill Tuco. The negative value Walter's measures

have for Tuco far outweighs their positive value regarding Walter and Jesse. So, all-things-considered, it seems Walter ought not to save Jesse from breaking his ankle. Nonetheless, when we consider just Jesse's interest in having an intact ankle, it seems right to say that he depends on Walter for the fulfilment of his interest. It even seems right to say that Walter has a strong presumptive reason to fulfil that interest of Jesse's. It is just that this presumptive reason is *defeated* by considerations that are *external* to the dependence relation. (I will briefly discuss such defeating considerations in §2.7.)

We then get the third necessary condition for dependency duties, to add to (1) and (2) above:

- (3) A's taking the measure in (2) would realise positive expected value regarding A and B.

### **2.5.2 Aggregative and Iterative Value**

It is still potentially unclear exactly which range of expected values are to be included in the assessment of a measure's value regarding agent and dependent. Should we just count the expected value of this one measure I'm considering taking now? Intuitively, that might seem right: if we're figuring out whether this particular measure is obligatory, then we want to know the value of *it*. But this creates problems. After all, small costs add up, both across time and at one time. Intuitively, we should sometimes be able to appeal to these aggregate costs as blocking the duty. For example, perhaps any one measure I might take to fulfil one person's important interests is not prohibitively costly for me, but I cannot take measures that fulfil *all*, or even *many*, persons' important interests without incurring prohibitive cost.

Suppose the aggregative, but not iterative, value is deemed non-positive by your theory—that is, the measures taken in aggregate have a non-positive expected value, but taken iteratively they each have a positive expected value.<sup>21</sup> Then, if each duty over each measure were dependent only on the expected value of *that one measure* being positive, then we would be unable to cite the overall non-positive expected value as excusing us from any one of the interest-fulfilling duties. But this cannot be right: I should be able to cite the aggregated expected costs of all these measures as disproportionate to their aggregated expected benefits, rendering the aggregate expected value non-positive.

---

<sup>21</sup> For example, suppose it takes you a day to save each of 20,000 lives. While it might be proportionate (i.e. have positive value) for you to spend one day saving one life, it might be that 20,000 days (i.e., almost 55 years) is too much for morality to demand of you, whatever the cause.

Garrett Cullity (2004a, Part II) gives a similar argument against iterative approaches, in particular as they apply to the principle “act beneficently.” Cullity notes that each potential beneficent act on its own is compatible with my possession of certain “life-enhancing goods” such as friendship, personal projects, and so on. But we cannot have these life-enhancing goods while acting beneficently in *all* circumstances. Cullity gives a complex argument, with which I will not engage here, for thinking that these life-enhancing goods should not be sacrificed at the altar of beneficence. If this is right, then we should take an aggregative, rather than iterative, approach to assessing whether (in my terminology) the value of one act of beneficence is positive. In short, we should consider the value of our *policy* of beneficence, not the value of individual beneficent acts, if we are to give due credit to the value of life-enhancing goods. Frances Kamm makes the same general point:

Even if there is no distinction between the cases as individuals, the cumulative effort or cost is different in aiding one or many. Even if there is no magic cutoff point such that the difference between aiding ten and aiding eleven, for example, will involve making more total effort than required, we can set an arbitrary cutoff so as not to go on aiding when the aggregate total would clearly be more than required. (Kamm 2000, 660; see similarly Hooker 1999)

If aggregative value can be non-positive even though iterative value is positive, this suggests that dependency duties are not constrained by the expected value of each interest-fulfilling measure considered in isolation, but by the aggregated expected value of all such measures. But exactly which measures’ value are we meant to be aggregating, here? Perhaps we should aggregate *all* measures that we could take to fulfil others’ interests. But dependency duties, in and of themselves, do not allow us to say this, because they do not even say that we have reason to take all these measures. Dependency duties arise out of being *best-placed* to fulfil interests. So it seems we should aggregate the value of measures that we could take to intentionally fulfil someone else’s important interest, where, in each case, we are best-placed to fulfil that interest (I’ll soon explain what “best-placed” means).

This gives us something like the following, as a necessary condition for a dependency duty:

- (4) If A were to take measures in all circumstances where (1)–(3) hold between A and any individual, then A would realise positive *aggregate* expected value regarding A and all those individuals.

At first glance, this seems like a plausible condition to add to conditions (1), (2), and (3) developed above.

### **2.5.3 A Ranking Method**

However, a problem remains. Suppose Peter is the only one capable of rescuing all of 100 children. Each of their interests is very important. It would realise positive expected value for him to rescue 10 of them, but non-positive expected value for him to rescue any more than that. (Imagine that rescuing 11 or more would cause him to develop severe hyperthermia, which serves as a “blocking” function on benefits.) If (4) is a necessary condition for a dependency duty, then Peter seems to be off the hook *altogether*: helping all 100 would realise negative value, so, by the lights of (4), he has no duty to help any of them.

This cannot be the right result. The right result in this case is that Peter has a duty to rescue of 10 of the 100. After all, that is the threshold where the expected value is still positive. Peter should be let off the hook for not helping all 100, but *not* off the hook for helping 10 of them. It seems that Peter should act up to the point where the expected value becomes negative. How, then, should Peter select *which* of the 100 he assists?

One possibility is to help them on a first-come-first-served basis, where Peter cites the cost of helping the first 10 when the 11th shows up. But this won’t work in cases where all 100 appear simultaneously. And method that deals with both simultaneous and queue-type cases would make for a more principled solution.

A plausible method is for Peter to formulate a ranking of the 100 measures, for saving each of the 100 children. Each measure is weighted according to the importance of the interest it is aimed at fulfilling, along with the other expected costs and benefits for Peter and for the relevant child. The weighted measures are then ranked. Peter’s duty is to work his way down the ranked measures, taking each until he reaches the limit of positive aggregate value. There might be several measures that are in a “tie” situation on the list: if Peter cannot take all the tied measures while realising positive value, then he is afforded discretion in choosing which of the tied measure(s) he takes.

Notice that on this weighting and ranking method, Peter’s duty is not “imperfect.” Roughly speaking, imperfect duties are those in which the agent is afforded significant

discretion in how to discharge the duty.<sup>22</sup> Given the strictness of the procedure I have just outlined for weighting and ranking interests, Peter is unlikely to enjoy *significant* discretion over how to discharge the duty (except in the case where numerous interests are tied, and fulfilling some but not all of those tied interests would realise positive value).

This weighting and ranking method is just one possibility. One's foundational theory might endorse a different method, in which Peter enjoys more discretion in fulfilling the interests of the 100, such that the duty is properly described as imperfect (most straightforwardly, if one's theory places great value for agents on having discretion). In that case, the weighting and ranking procedure itself would realise negative value for the agent, and perhaps negative value for agent and dependent combined. But this does not arise from within the internal logic of dependency duties, which is what I am here exploring. The ranking method, I suggest, allows us to address the concern that aggregate (but not iterative) cost to the agent might be prohibitively high, while still insisting that agents do what they can for as many dependents as they can within reasonable cost.

Condition (4) above then becomes disjunctive:

(4) Either: if A were to take measures in all circumstances where (1)–(3) hold between A and any individual, then A would realise positive *aggregate* expected value regarding A and all those individuals;

Or: when the importance in (1) and value in (3) are used to weight the measures in (2), these measures rank sufficiently highly among similarly weighted measures (for which (1)–(3) also hold), such that the measures in (2), *and* all more highly-ranked measures, could be taken with positive aggregate expected value (regarding A and those whose important interests are thereby targeted).

#### **2.5.4 Packaging Interests**

An inverse case is one in which aggregation *renders* the expected value positive, due to something like economies of scale. What if Peter can assist *all* of the 100 at positive aggregate value if and only if he “packages” their interests? For example, it might be far too costly for Peter to jump in and save all of 100 drowning children, as he would die of

---

<sup>22</sup> I say “significant” because, as Roderick Chisholm (1963) points out, all duties involve *some* latitude. Just how much discretion is “significant” is open for debate, but there are some clear cases on either side of the line: having discretion to rescue a drowning child with one's left hand or one's right hand is not sufficiently significant to render the rescue duty imperfect; having discretion about which one of 100 charities to give to does render the charity duty imperfect. On distinguishing perfect from imperfect duties, see: Igneski 2006; Rainbolt 2000.

hyperthermia. He could jump in after some of them, but jumping in after all of them would realise negative value. He's the only one around. What if he considers their interests as a package? For example, what if he can *predict* that he will *soon* face 100 drowning children (though there are none around now), and he considers, rather than jumping in after each one to fulfil their interest in not-drowning, simply building a fence around the pond? Suppose that building a fence would be, for Peter, less costly than jumping in 100 times—and would realise positive expected value regarding himself and all the dependents, considering all the important interests he would fulfil. Intuitively, he should build the fence. More generally, the intuitive pull behind dependency duties seems to dictate that we should package interests in a way that allows us to take the most efficacious measure for fulfilling all of them.

There are a number of ways to package important interests. We can package the interests of one individual over time; of one individual at the same time; of many individuals over a stretch of time; of many individuals at the same time. By packaging interests in any of these four ways, we will get a different target *set of* interests, and the value of the measures various agents can take to fulfil them will vary. The value of these measures will be a function of costs and benefits to the agent, combined with costs and benefits to *all* those dependents whose interests are in the set.

We can make this explicit by saying the following: when determining whether A has a dependency duty to take some measure to fulfil B's important interest, we should consider the most efficacious measure open to A. But this measure is not just the measure most likely to fulfil B's interest (as we have already stipulated). Rather, if a number of A's measures are *tied* in respect of what is most likely to fulfil B's interest, then the most efficacious measure is the one of those tied measures that best fulfils *other* important interests, whether of other dependents or of the same dependent. Then *all* these other interests can contribute to the measure's value, since the agent stands in a dependence relation to the bearer of each of them. Thus if Peter has a choice between either jumping in after Jonny, or building a fence to protect Jonny and other children (where these are both equally likely to save Jonny), he should build a fence because that is the measure that—among the measures that are tied for likelihood of success regarding Jonny's interest—realises most value with regard to Peter and *all* dependents whose important interests are up for fulfilment by the measure.

This somewhat technical consideration does not require that we add a new condition to our growing set of necessary conditions for dependency duties. But it is worth bearing the idea of “packaging interests” mind when we are considering which measure is “most efficacious” for fulfilling some important interest. If a number of measures are

equally likely to fulfil an interest, then the most efficacious measure is the one that realises the most expected value with respect to *other* interests that might be fulfilled by that very same measure. The idea of using one measure to fulfil a large number of important interests will re-emerge in Part III's discussion of R2P.

### **2.5.5 The Principle So Far**

We can now say that some of the necessary conditions for A's incurring a perfect dependency duty to take measures to fulfil B's interest are:

- (1) B has an important final interest, FI, that is unfulfilled, where the interest is described in a way that does not include morally irrelevant facts (though it might be temporally or otherwise indexed); and
- (2) If: A intentionally takes measure M, where M is the most efficacious measure open to A to fulfil the interest in (1); then: even if other agents do not coordinate with A, FI will be fulfilled with a likelihood that is *proportionate* to FI's importance, where more important interests are proportionate to a wider range of likelihoods (with a lower bound that is lower); and
- (3) A's taking the measure in (2) would realise positive expected value regarding A and B; and
- (4) Either: if A were to take measures in all circumstances where (1)–(3) hold between A and any individual, then A would realise positive *aggregate* expected value regarding A and all those individuals;  
Or: when the importance in (1) and value in (3) are used to weight the measures in (2), these measures rank sufficiently highly among similarly weighted measures (for which (1)–(3) also hold), such that the measures in (2), *and* all more highly-ranked measures, could be taken with positive aggregate expected value regarding A and those whose important interests are thereby targeted.

This does not capture the possibility of imperfect (discretionary) duties, as would arise when there is a tie in the rankings of weighted measures. The tie possibility can be captured by saying the following, which will become part of the consequent, rather than the antecedent, of the principle for dependency duties:

If several measures are equally ranked by the methods in (4), and if A could take *some but not all* of the equally-weighted measures while realising positive aggregate

expected value, then A has a duty to take some of the tied measures, up to the threshold of positive aggregate expected value.

## 2.6 “Least Costly”

Conditions (1)–(4) do not get us to dependency duties, understood as duties that accrue to the agent *best-placed* to fulfil some important interest. If conditions (1)–(4) were the end of the story, then multiple agents could bear duties to fulfil exactly the same interest, if they all met (1)–(4). It is possible that there are numerous agents who are sufficiently capable of fulfilling an important interest at not-too-high iterative and aggregative cost to themselves and dependents. In fact, (1)–(4) might be a good set of individually necessary and jointly sufficient conditions for duties of mere capacity. Such a proliferation of duties would be an acceptable result, as we could say that each agent who meets (1)–(4) has a duty to take their respective measures, *until* one of them does so (or clearly will do so), at which time the others lose their duties (while, perhaps, acquiring new ones to make sure that the actor succeeds).

But I am interested in picking out one from among the many agents who have duties of mere capacity. This agent will have a stronger duty than all the other merely capable agents, and it will be this agent upon whom we will want to call first in cases where it would be disastrous if all agents with a duty of mere capacity acted. Imagine a crowd of capable people all jumping in to save the drowning child, getting in each other’s way and causing further drownings. We don’t want this. We want a single agent to act. Dependency duties pick out this one agent. This agent has a stronger duty than all the other agents, even if it wouldn’t be disastrous for all to act. And if it *would* be disastrous for all to act, then she has the duty in the first instance.

How, then, should we pick out the *best-placed* agent, so that we can know who has the dependency duty? The agent best-placed to fulfil some interest is the agent who meets conditions (1)–(4) and condition (5). A first pass of condition (5) is:

(5) A is the agent whose measure (from (2)) is such that, if A takes that measure, A would realise no less expected value (regarding agent and dependent) than would any other agent who meets (1)–(4) if they took their measure.

Note that this condition is importantly different from “A is most capable.” The most capable agent is simply the one whose most efficacious measure is *no less likely to fulfil* the relevant interest than is any other agent’s most efficacious measure. But the most capable

agent might incur great costs—or place great costs on the dependent—if she takes this measure. For this reason, we should not be concerned merely with the likelihood of her measure’s success, but with the other costs and benefits of her measure for her and the dependent. The likelihood of her measure’s success at fulfilling the interest definitely *contributes* to its expected value, but it is not the only factor involved.

Notice also the condition (5) does not render conditions (2)–(4) redundant. Condition (5) does not presume that *any* agents, including A, actually meet conditions (2)–(4). Thus, A might meet condition (5), but be insufficiently likely to fulfil the interest in (1), even with her most efficacious measure. Then she would not meet condition (2). Or A might meet condition (5), but her most efficacious measures would realise too high a cost regarding herself and the dependent. Then she would not meet condition (3). Or A might meet condition (5), but be unable to take measures in all cases where someone depends on her—at least, unable to do so at not-too-high a cost. Then she would not meet condition (4). In none of these cases do we want to say that A is best-placed to fulfil the important interest. Thus (2)–(4) are necessary conditions for a dependency duty, and (5) does not supplant them.

But the first pass above is not quite the end of the story for “least costly.” In characterising “capacity,” I said that other agents are treated like any other factor in the agent’s environment (except in cases where they intentionally coordinate with you to fulfil the interest, to be dealt with in Chapter Three). This matters, because the “best-placed” agent—according to the above first pass at condition (5)—might not do what her moral duty demands. Suppose Ash has broken his leg, and his housemate, Brandy, is best-placed (according to condition (5)’s first pass) to take Ash to the hospital. Yet Brandy is extremely lazy and callous, and will not bother doing her duty. For the purpose of determining *Brandy*’s duty, we do not allow her unwillingness to take Ash to hospital to impact on her duty to do so.

But what about the duties of Cara, Ash and Brandy’s next-door neighbour, who knows Brandy will selfishly watch television while Ash suffers, and that there is nothing anyone can do to convince Brandy otherwise? What is Cara’s duty? She does not meet condition (5)’s first pass: Cara’s taking Cara’s measure would realise less value than Brandy’s taking Brandy’s measure. But, given Brandy’s laziness, Ash surely depends on Cara. Let us consider why this seems right.

Intuitively, this seems right because in figuring out whether some agent has a dependency duty, we should not assume that others will comply with morality. We should be as realistic as we can about the expected probability that others will comply with their duties, just as we are realistic about natural events that might intervene and disrupt the

attempts of the agent under consideration. By contrast, when figuring out whether some agent has a dependency duty, it is obviously not relevant whether *she* will or won't comply with morality. Agents should not be able to get out of assisting others simply by not wanting to. When determining Brandy's duty, then, we treat her as an agent, who cannot get out of her duty by simply not doing it. So she is best-placed, on that assumption. This gives Brandy a dependency duty.

When determining Cara's duty, however, we treat Brandy as a feature of the environment to be worked around. Given that Brandy will not do her duty, Cara is best-placed. Because we assume that the duty-bearer *will* comply, but are realistic about whether others will comply, it seems that there are sometimes two (or more) agents with dependency duties to take different measures to fulfil the same interest. This is simply because, when we are considering the duties of the two different agents, we hold different things fixed. Of course, if it's possible they might be convinced or coerced, then that is another matter (to be addressed in Chapter Three). For now, assume Brandy and Cara have no means of interaction.

So, *given Brandy will not do as morality says she should*, Ash depends on Cara for his trip to the hospital. Cara is second-placed in absolute terms, but she is best-placed given the non-compliance of others. And we take this non-compliance as a given fact when determining the value of Cara's taking certain measures: if Cara does not take the relevant measures, Ash will continue to suffer. Thus Brandy has a dependency duty, *and* (given that Brandy will not comply) Cara has a dependency duty. The upshot is this. When we are considering an agent's duty, what matters is not that the agent's measures are not worse than anyone else's, but that they are not worse than those of any other agent *who will do their duty if they have one*.

We can thus slightly refine condition (5):

(5) Of all agents that meet (2)–(4), A's measure (from (2)) is such that, if she were to take that measure, she would realise no less expected value (regarding agent and dependent) than the expected value (regarding agent and dependent) that would be realised by any other agent *who would take measures to fulfil this interest if they had a duty to do so*.

The italicised clause has the implication that I do not have to give an account of “back-up” duties—that is, of duties that accrue to the second best-placed agent when the best-placed agent reneges on his duty. If the best-placed agent is not willing, then the second best-placed agent will meet the condition of “best-placed” given above. This is

because, for each of them, we are considering a different set of agents and are saying that he is the best-placed agent *within that set*. This set is: the agent in question plus all other *willing* agents. If the agent is not willing, then he will not be included in others' sets. This means he cannot be best-placed of those sets, and that someone different will be (if anyone is). But he will always be included in his own set, whether or not he is willing.

If several agents are “best-placed within his or her own set,” then there are multiple “best-placed” agents—even if one of them is (in some colloquial sense) “objectively” better-placed, in that he is the best-placed if we take *all agents (whether or not they are willing)* as the relevant set. But this is not the relevant set, because one agent’s unwillingness should make a difference to whether or not other agents have duties, but should not make a difference to whether or not he himself has a duty. (There is an important question about whether the “objectively” best-placed agent has a duty to compensate those who actually end up fulfilling the interest. The answer to that question requires drawing on considerations that are external to the dependence relation, and which I lack space to address here.)

Notice as well that condition (5) holds that the measure “would realise no less expected value.” This is different from saying that the measure “would realise more expected value.” The upshot is that, if two agents would realise equal value, then they both have a duty to take the measures. But once one of them starts to take the measures, the other’s duty is likely to be dissolved. This is because, once the first agent has started acting, it is unlikely that the second agent would realise as much expected value by acting as the first agent would by continuing the act. The reason why we should allow both agents to count as “best-placed” in the first instance is this: if we did not allow them both to count as best-placed, then neither would have a dependency duty. They might both have duties of mere capacity, but these would not necessarily be stronger than the capacity-based duties of any other merely capable agent. This is a bad result: we want one of these two agents to act, not just any other agent who is less well-placed. So, we should allow that both are (equally) best-placed. The above condition allows this.

## 2.7 The Dependency Principle

The forgoing discussion has led us to the following principle:

If:

- (1) B has an important final interest, FI, that is unfulfilled, where the interest is described in a way that does not include morally irrelevant facts (though it might be temporally or otherwise indexed); and
- (2) If: A intentionally takes measure M, where M is the most efficacious measure open to A to fulfil the interest in (1); then: even if other agents do not coordinate with A, FI will be fulfilled with a likelihood that is *proportionate* to FI's importance, where more important interests are proportionate to a wider range of likelihoods (with a lower bound that is lower); and
- (3) A's taking the measure in (2) would realise positive expected value regarding A and B; and
- (4) Either: if A were to take measures in all circumstances where (1)–(3) hold between A and any individual, then A would realise positive *aggregate* expected value regarding A and all those individuals;  
Or: when the importance in (1) and value in (3) are used to weight the measures in (2), these measures rank sufficiently highly among similarly weighted measures (for which (1)–(3) also hold), such that the measures in (2), *and* all more highly-ranked measures, could be taken with positive aggregate expected value (regarding A and those whose important interests are thereby targeted); and
- (5) Of all agents that meet (2)–(4), A's measure (from (2)) is such that, if she were to take that measure, she would realise no less expected value (regarding agent and dependent) than the expected value (regarding agent and dependent) that would be realised by any other agent who would take measures to fulfil this interest if they had a duty to do so;

Then (6), (7), and (8):

- (6) In the absence of strong defeaters, A has a perfect, all-things-considered duty to intentionally take M; and
- (7) If the ranking in the second disjunct of (4) ranks multiple measures equally, and if A could take some but not all of those equally-weighted measures and realise positive aggregate expected value (regarding A and dependents), then, in the absence of strong defeaters, A has a duty to take some of the tied measures, up to the threshold of positive aggregate expected value (regarding A and dependents); and
- (8) The duties in (6) and (7) are stronger—A does more moral wrong if he defaults on them—if the interest in (1) is more important and the value in (3) is higher.

The Dependency Principle is a conditional: if (1)–(5), then (6)–(8). However, I noted earlier that there is an important supplement to the Dependency Principle. This is not part of the principle itself, and is an unconditional imperative:

- (9) All agents have a duty to not intentionally diminish their capacities primarily in order to avoid dependency duties that might arise.

In Parts II and III, I will regularly have reason to draw upon the details of the Dependency Principle. These details will sometimes be crucial for resolving ambiguities in, and making distinctions regarding, the doctrines of care ethics and R2P. But the reader is not expected to keep all these complexities in mind: I will explicitly state the relevant condition(s) as necessary. The most important parts of the principle are the following: measure M is sufficiently likely to fulfil an important interest; M would realise positive expected value (regarding the agent and dependent) if taken in this instance; positive expected value would be realised (regarding the agent and dependents) if like measures were taken in like circumstances; and M would realise no less expected value (regarding the agent and dependent) than the most efficacious measure of any other agent who would discharge a duty if they had one.

An even more “bare bones” version of the principle is just this. A has a dependency duty to take measure M when A’s most efficacious measure for fulfilling an important interest: is *sufficiently likely* to fulfil the interest; would realise *positive expected value* regarding agent and dependent; and would realise *no less expected value* regarding agent and dependent than any other willing agent’s most efficacious measure for fulfilling the interest. This simplified version elides many of the details, but will be sufficient for some (though not all) of the uses of the Dependency Principle in Parts II and III.

The antecedent of this conditional is not *necessary* for a duty to take measures that fulfil someone’s important interests. Such duties might arise out of any number of other principles—for example, if the agent has promised to do so, has caused the non-fulfilment of the important interests, or similar. But I will assume that the antecedent is necessary for the duty to be properly called a “dependency duty.” Up until this point, I have been talking about “dependency duties” loosely. Having arrived at the preceding account of those duties’ conditions, from now on I will be concerned with those duties as they arise under these conditions. From now on, I will call the duties that arise out of the Dependency Principle “dependency duties.”

An attractive feature of this principle is its flexibility. It allows us to plug in whatever conception we like of what interests are important, what makes them important, and how various benefits and costs are to be weighed up. It is a framework that maps how various concepts (important interest, benefit, cost, etc.) interact so as to produce the dependency duties that are common to a range of ethical and political outlooks. On different conceptions of these, dependency duties will arise in somewhat different circumstances. This is as it should be, since it is plausible that dependency duties exist, but foundational theories will disagree on precisely when those duties exist. The principle is intended to capture the consensus, not settle the disagreements.

As consequents (6) and (7) state, the Dependency Principle produces all-things-considered duties only in the absence of strong defeaters. These defeaters might arise from other moral principles, or perhaps from non-moral principles and values. Conditions (1)–(5) capture what the dependence relation contributes to the all-things-considered “ought,” but this relation is not the only contributor to that ought. There might be very strong moral reasons for taking measures to fulfil an important interest, even though other moral reasons come into play such that the duty in (6) or in (7) is defeated.<sup>23</sup> In cases with defeaters, we can, in line with the spirit of the Dependency Principle, acknowledge that a strong *pro tanto* dependency reason exists. Indeed, this *pro tanto* reason is the driving force behind dependency duties, even if it is pushed back by strong defeaters in some cases, to foreclose the possibility of an all-things-considered duty.

Recall the case of Walter. Walter had a very good chance of rescuing Jesse at low cost to himself and Jesse, but would kill Tuco in the process. Now imagine that Walter is joined by a second agent, Skylar, who has a slightly lower chance of rescuing Jesse at slightly higher cost to herself and Jesse, but no third parties will be harmed if Skylar does it rather than Walter. It is implausible to assert that Walter has an *all-things-considered* duty and that Skylar does not. Nonetheless, if we focus just upon Jesse, and ask who is best-placed to rescue him, it seems right to say that Walter is best-placed. Jesse depends on Walter. Walter is better-placed than Skylar is *with regard to Jesse’s important interests*; he is just not better-placed all-things-considered. Walter has a strong presumptive reason to rescue Jesse, even though there is a strong defeating consideration against his doing so and in favour of Skylar doing so—namely, the death of Tuco. Walter has a *pro tanto* dependency duty, but not an all-things-considered one.

---

<sup>23</sup> For different conceptions of the moral and other values that can outweigh or constrain duties in general, see Wolf 1982; Chappell (ed.) 2009. On interactions between different duty bases, see D. Miller 2001.

There are numerous important debates here, with which a general framework for dependency duties need not engage closely. There are, for example, the debates on the potential defeating effects of fairness, distance, nationality, and association.<sup>24</sup> Philosophers writing on these topics generally agree that we have duties to do something (however minimal) for those whose most important interests are unfulfilled and whose lives we can greatly affect. (Though there is not usually a distinction between whether we are *best-placed* to do so—and thus whether these people are *dependent* on us, in my sense—or whether we are merely *able* to do so.) These philosophers address the question of how these “assistance” duties weigh up against duties to our nearest and dearest, to co-nationals, or to those nearby (who may not be co-nationals). My account of dependency duties is neutral on this question. On my account, such values should be factored into (6) and (7) as potential strong defeaters—but exactly how they factor here is left open. Suffice it to say that conditions (1)–(5) generate a strong presumptive duty, which only a weighty countervailing consideration could defeat.

## 2.8 Conclusion

Having developed my preferred principle for when dependency duties exist, I have already gone some way to answering the question of their scope. For example, since important interests and sufficient capacities are both graded concepts, and since the duty depends on the importance of the interest being *proportionate* to the strength of the capacity (i.e., the likelihood of success given measures), there can be dependency duties to fulfil some very important interest even if the capacity is quite weak. There can also be duties to fulfil some not-very-important interest if the capacity is strong enough. Yet dependency duties are circumscribed by cost to the agent and the dependent; and they might be defeated by other moral and non-moral values. Depending on what other values there are, dependency duties may not be as demanding as theorists like Singer make them out to be. Yet, as we shall see in Parts II and III, my way of precisifying the basic thought gives us many interesting, substantive, and real-world normative claims.

Before considering these real-world uses of the Dependency Principle, however, it is important to delve more deeply into the implications of the basic thought with which I began. This thought—if you are dependent on someone in the right way, then they have a moral duty—seems to apply whether the “someone” is a group or an individual. In fact,

---

<sup>24</sup> On fairness, Murphy 1993, Murphy 2000; on distance, Chatterjee (ed.) 2004; on nationality, Miller 2007; on association, Scheffler 1997; Scheffler 2001.

Part III of this thesis will be entirely concerned with group applications. I have touched upon this already, when noting that the child and the father cannot each have a dependency duty to make the cake, because this would ignore the importance of the other's intentional cooperation. They can only have a duty to do so together. I will now consider how these kinds of duties—duties where some *group* of agents seems best-placed to do something—should be best understood.

## Chapter Three:

### Collectives' Duties and Coordination Duties

#### 3.1 Introduction

It is impossible to do justice to the scope of dependence-based duties without discussing the duties of groups—whether those groups are highly organised, extremely ad hoc, or somewhere in between. To see this, simply consider the myriad circumstances in which a *group* seems best-placed—in the sense I developed in Chapter Two—to fulfil some person’s important interest. This is most clear in the case of organised groups. Take Alma’s interest in having her good health socially guaranteed against standard threats. If Alma lives in a state with a public health system, then the state (or the government, or the department of health) is likely best-placed to fulfil this interest. Or take children’s several interests in having a fence built around a pond. Depending on who has control over the land around the pond, the agent best-placed to fulfil those interests might be the local council.

Being best-placed (in the sense I described in the previous chapter) entails, *inter alia*, having measures open to you that are sufficiently likely to fulfil an important interest. This is an important part of the duty’s ground. Since agents can take only their own measures, it would be quite odd if agents could incur dependency duties to take the measures of some agent to which they are not identical. If you have measures available to you for rescuing a child, but I don’t, then it would be very strange if your measures could somehow generate a duty for me. If this is right, then, in order for a *group*’s measures to generate a duty, the measures must generate a duty *for the group*: the group has the measures, so only the group can take them, so only the group can have a duty to take them.

But groups are collections of individuals, so surely if a group has a duty to take some measures, that must have implications for the group’s members. Surely groups can’t do anything over and above what their members can do together. So the group’s duty must imply a number of individual duties—but individual duties that are somehow grounded in the *group*’s being best-placed to fulfil some interest, because the measures belong to the group. How are we to make sense of the nature of groups’ measures and duties, and exactly how they generate duties for individuals?

To answer this question, I begin (in §3.2) by providing a typology of groups and arguing that only one type—which I will call *collectives*—have the agency necessary to bear duties, including dependency duties. This conclusion is worth arguing for in depth, as a

surprisingly large number of groups to which theorists assign prospective or retrospective responsibilities—such as those whose members have shared interests,<sup>25</sup> shared attitudes,<sup>26</sup> similar resources,<sup>27</sup> or a shared culture<sup>28</sup>—seem not to be collectives in my sense. If I'm correct about the conditions a group must meet to have the moral agency necessary to bear duties, then we must restrict the scope of groups' duties—including groups' dependency duties—accordingly.

Having established which groups can have duties, §3.3 addresses a set of problem cases. These are cases where there is an important interest that is not fulfilled, but no existent individual or collective agent is sufficiently well-placed to fulfil the interest to have a dependency duty to do so—say, no existing agent's measures are sufficiently likely to fulfil the interest, or no existing agent's measure would realise positive expected value for agent and dependent. Intuitively, this is not necessarily the end of the story for dependence-based duties. Imagine that, if a number of individuals were to *form* a collective, or were to *reform* an existing collective, then that formed or reformed collective would incur a duty according to the Dependency Principle. Imagine, moreover, that the individuals could each take steps towards forming (or reforming) such a collective while realising positive expected value regarding themselves and the dependent. Alternatively, imagine that an important interest would be fulfilled if and only if a number of individual agents acted *responsively* to one another, though without constituting a collective agent, and that each individual could take responsive actions while realising positive expected value regarding themselves and the dependent.

§3.3 argues that, in the vast majority of these kinds of cases, none of the individuals can have dependency duties (or any other kind of duty) to form or transform a collective, or to act responsively to one another. This is because each individual has no measure available to him or her—even diachronically—that is sufficiently likely on its own to form or transform the collective, and none of them can—even diachronically—unilaterally act responsively in a way that fulfils an important interest. So, none of the individuals can have a dependency duty regarding the interest. Yet neither can *the group* have a duty to fulfil the

---

<sup>25</sup> Joel Feinberg (1968, 687) describes a train robbery in which passengers could save the day by acting together. He concludes that “the group was collectively but not distributively at fault.” See similarly Held (1970, 477).

<sup>26</sup> Larry May and Robert Strikwerda (1994) argue that men are collectively responsible for rape. This is distinguishable from the claim that *individuals* are responsible for harms perpetrated by their attitudinal communities, to the extent they share those attitudes. On the latter, see May (1992).

<sup>27</sup> Cullity (2004a) argues that the affluent are “collectively” morally required to contribute to international aid agencies.

<sup>28</sup> D. Miller (2004) argues that nations (cultural groups distinct from states) can be responsible for outcomes.

interest since, by hypothesis, this group is either not an agent or does not have any interest-fulfilling measure open to it (not even diachronically).

Intuitively, individuals in these cases do have duties to form or reform the group, or to act responsively to other individuals. To capture the moral imperatives in such situations, we need to posit two new kinds of duties for individual agents: first, duties to take steps to *form* or *reform* collectives so that those collectives can bear, and fulfil, dependency duties; and second, duties to act *responsively* to other individuals with a view to bringing it about that an important interest is fulfilled.

These two kinds of duties arise in different kinds of situations, as we shall see. Both are what I will call “coordination duties.” Coordination duties arise out of the non-fulfilment of important interests and out of agents being “best-placed” to fulfil those interests in some very loose (and technically false) sense. Thus coordination duties are very close cousins of dependency duties. However, I will reserve the phrase “dependency duties” for those duties that arise out of the Dependency Principle. I will spell out coordination duties in the Coordination Principle, which serves as an important companion to the Dependency Principle.

§3.4 discusses further issues for the duties of collective agents. It examines what it means for a collective agent to have duties, including dependency duties. I argue that collectives’ duties reduce to individual duties, though the organisational structure of a collective is necessary for individuals to fulfil the duties to which collectives’ duties reduce. This analysis of collectives’ duties is not particular to dependency duties—it applies equally to collectives’ duties based on having collectively caused harm, collectively making a contract, or so on. But given the large role that collective agents play in real-life dependence-based duties (as shall be demonstrated in Part III), a deep understanding of the nature of collectives’ duties is important for understanding the nature and scope of dependence-based duties.

## 3.2 Duty-bearing Groups

### 3.2.1 Agency

I will assume that each duty is the duty of some moral agent—at least, if that duty is a duty to act or refrain from acting, as dependency duties generally are.<sup>29</sup> This assumption is based

---

<sup>29</sup> I say “generally” because we shall see in Chapters Four and Five that there can be dependency duties to have certain attitudes. While the text only discusses actions, I assume these remarks apply equally to attitudes. I discuss attitudinal duties in §4.5.4.

on the natural thought that only agents can act (or refrain from acting—for ease of exposition, I will talk only about acting, rather than refraining. I will assume that they require the same cognitive equipment—indeed, refraining could be a kind of acting). If only agents can act, then only agents can have duties to do so. So, if a group can act, then the group is an agent. Then and only then can it bear duties.<sup>30</sup> Thus if we want to figure out which groups can bear duties to act, a good place to start is with thinking about which groups can act.

On a basic functionalist “belief/desire” model of acting, acting requires having (1) something that plays the role of reasons (e.g., desires, goals, preferences), which, in combination with (2) something that plays the role of beliefs about one’s environment, move one to make (3) something like decisions about how to act in that environment.<sup>31</sup> Let us assume that agents make decisions by processing reasons and beliefs regarding them, that is, by implementing a decision-making procedure. Decision-making procedures are procedures that move from (1) and (2) to (3). Of course, agents might make decisions in other ways—but we’re concerned with those decisions agents make when *moral reasons* (specifically, duties) play the (1) role. These decisions presumably require a decision-making procedure. So let us assume that each action-demanding duty belongs to a moral agent: to an entity with a decision-making procedure that can process moral reasons to produce decisions on which the entity acts.

We often have more than this in mind when we speak of agents: complex attitudes, phenomenology, conceptual understandings, and so on. Groups are obviously incapable of such full-fledged human agency. But perhaps they are capable of at least approximating the minimal conditions for moral agency, by having group-level decision-making procedures that can take in and process reasons (including moral reasons), combine these with beliefs about the environment, produce decision-like outputs where the decisions is “to perform action  $\varphi$ ,” and then performing the actions they have decided to perform.<sup>32</sup> Which kinds of

---

<sup>30</sup> Bill Wringe (2010) argues that non-agent groups—specifically the “global collective” that is the conjunction of all agents in the world—can be the “subject” of duties, though such agents cannot be the “addressees” of duties. The addressees, for Wringe, are those whose “behaviour or attitudes might be affected if they came to accept” that the duty exists. He accepts that only agents can be addressees. But just what it means to be a “subject” on his view is left quite unclear. (It seems on his view that the table might be the subject of my duty to lift the table, for example. But if this is right, then it is not clear why it is at all normatively interesting to say that non-agent groups can be subjects of obligations.) In any case, when I say that an agent “bears” a duty, I mean that they are the “addressee” of the duty, in Wringe’s terms.

<sup>31</sup> Along these lines, Christian List and Philip Pettit (2011, 20) give three features of an agent: it has representational states (beliefs), motivational states (aims), and the capacity to process these states such that “in favourable conditions, within feasible limits, it acts for the satisfaction of its motivations according to its representations” (2011, 20).

<sup>32</sup> If one thinks acting requires having intentions, and rejects collective intentions from the outset, then bear in mind that the intentional component of a collective’s action can be reduced to the members’ several

groups might be capable of this? Answering this question will require detailed analysis. But this will be worth it: it will allow us to correctly conceptualise the dependence-based duties that seem so often to belong to groups, and to know exactly when a group can have these duties and when they cannot.

### ***3.2.2 Random Aggregates***

We can begin by ruling out some groups as clearly not moral agents, and therefore clearly incapable of incurring duties, whether dependency or otherwise. I will call these “random aggregates.” These are conjunctions of individuals, where each individual in the conjunction makes decisions (including forming intentions to act) completely independently of the others. Any random conjunction of individuals can be included in this category. These combinations can be completely without significance (for example, “me, William Shakespeare, and the old woman in that café in France”), though the members might be identified by reference to a common characteristic (“the beach bathers,” “the restaurant diners”).

The crucial characteristic of a random aggregate is this: none of the members are *responsive* to one another, and they do not constitute a *collective* (terms that are defined below). Because this category is so broad, there is little homogeneity among random aggregates. This means that some random aggregates are more useful for moral philosophy than others. The moral usefulness of a random aggregate is largely determined by the ease with which its members can become responsive, or come to constitute a collective, to bring about a morally desirable action or outcome.<sup>33</sup> This ease can often be tracked by properties such as the proximity of members to each other and the ease of communication among members. For this reason, “the beach bathers” is probably a more useful random aggregate than “me, Shakespeare, and the old women in that café in France”—but they are both random aggregates nonetheless.

### ***3.2.3 Responsive Individuals***

Moving up a level in terms of usefulness, we find aggregates of responsive individuals. Consider a group of bathers on the beach.<sup>34</sup> Suppose they did not come to the beach

---

intentions to play their part in achieving the collective’s action (Tuomela 2006, 36–7). I discuss related issues in §3.4.2.

<sup>33</sup> I say “largely” because random aggregates can be extremely normatively significant as moral patients—as when a particular population is the target of genocide or terrorism. (Part III will focus on cases of this kind.) In the text, I am thinking of random aggregates’ uses for moral philosophy as *agents*.

<sup>34</sup> My discussion of this example in this section and the next builds upon that in Collins 2013.

together. Once at the beach, they might (though they will not necessarily) be minimally responsive to each other. For example, suppose Ben infers that Laura prefers that other beachgoers not put their towels right next to Laura's. Ben might infer this from the way Laura scowls at him, or maybe just because Ben and Laura both live in a society in which the majority of people have this preference.

Ben is responsive to Laura insofar as he puts his towel at a distance from hers *at least partly because* he believes it is her preference that he does so. (Laura's preference need not be the only reason, or even the decisive reason, for Ben's placing his towel where he does.) At this point, Ben and Laura cannot use their responsiveness to bring about any outcome other than keeping out of each other's way. The same is true if there is a total of five, ten, or a hundred people on the beach being responsive in this minimal way.

As I will use the term, an individual, A, is “responsive” to another, B, just in case A's beliefs about B's decision-making procedure's beliefs and desires (inputs) and decisions (outputs) affect the decisions of A's procedure in a way that responds positively to what A perceives to be B's (objective or subjective) reasons. This might involve A's acting in a way that B does not desire, if A chooses to respond to B's objective reasons and A believes that B's desires do not properly track B's objective reasons.

Crucially, A's responding positively to B's perceived reasons could amount to A's responding to B with a view to producing some morally desirable outcome that A believes B has reason to have a hand in producing. This might take the form of *working together* with B in order to produce the outcome (through *mutual* responsiveness). It might take the form of A acting responsively to B so that, at a later time, B becomes capable of producing the outcome *on B's own* (in this case, A might act unilaterally responsively to B). Or it might involve acting *manipulatively* toward B to evoke a morally good outcome that B would not have produced except in response to A's intervention. Or it might involve acting *adaptively* toward B, such that A takes up the slack that she realizes will arise from B's failure to act entirely as he ought.

To see this, consider some examples from Chapter Two. If a father believes that, say, his six-year-old child has good reason to have a hand in baking a cake, then the father might act responsively to the child by *working together* with the child with a view to the cake's being baked. Working together requires mutual responsiveness: it requires that the child is responsive to the father with a view to baking the cake, as well as vice versa.<sup>35</sup> This is a multilateral case; a case where agents are multilaterally responsive.

---

<sup>35</sup> This might sound overly demanding of the child—but recall from the sunbathing case that minimal responsiveness can be as much as placing your towel a certain distance from someone else's. The child will

As a unilateral case, suppose Peter's doctor cures Peter's chronic and debilitating fear of water, and suppose the doctor does this in response to the reason she believes Peter has to be such that he can save drowning children. In this case, the doctor is acting responsively to Peter, but it might be false that Peter is acting responsively to the doctor—false that Peter is responding positively to (what Peter perceives to be) the doctor's reasons, beliefs, and decisions. Then, the responsiveness is unilateral rather than multilateral.

The same goes for another case from Chapter Two: third parties' actions (and possible duties) to make my parents such that my parents can love me. These are actions of being responsive to others, by acting on those others, possibly unilaterally. Any *duties* that we have to act on others in this way—for example, duties to make others such that they can (on their own) fulfil some important interest—will be duties to be *responsive* to those others, for example by responding positively to the reasons the other has to be such that they can fulfil some important interest.

In the limiting case, responsiveness is unilateral and not commonly known. But agents can be more or less responsive depending on the extent to which (and number of reasons on the basis of which) they respond positively to one another. The extent to which a collection of individuals is responsive depends on, *inter alia*, how many individuals are responsive to how many others, how many decisions of each individual are affected by the others' perceived reasons, and how strongly others' perceived reasons influence the decisions of responsive individuals.

Good outcomes can result from responsiveness. Imagine there are six strangers at the beach: one drowning and the others sunbathing. Suppose each sunbather infers that every sunbather desires that the swimmer be rescued, and believes this gives each a reason to have a hand in this outcome. The swimmer can be reached only with a motorboat. It will take two people to drag the boat to the water and hold it while a third starts it. The boat will take off straight away, so the fourth and fifth, who will pull the swimmer into the boat, must already be in the boat. Two are required to pull in the swimmer because none is strong enough to do so alone.

Suppose one beachgoer, Ben, starts trying to drag the boat. Another, Laura, sees him and decides to help because she believes it will satisfy each of their reasons to have a hand in effecting the rescue. Two others, Jon and Julie, run to the water, ready to jump in the boat once it is there. They do this because the others are already dragging the boat, otherwise they would drag themselves. Stan sees what the others are doing, remembers he

---

presumably at least believe the father has decided to aim at a state of affairs in which the cake is baked. If the child acts responsively to this perceived decision, the child is being responsive.

is good at starting motor boats, and so runs to the water to be ready to start it. Jon and Julie each jump in the boat once it gets there because they can see Ben and Laura are now tired. Ben and Laura each hold one side of the boat because each believes this will help Stan start the boat. Stan starts it. When the boat arrives at the swimmer, Jon and Julie each grab one of the swimmer's arms and pull. The swimmer ends up in the boat.

No beachgoer could have rescued the swimmer alone. Moreover, the complementary nature of the individuals' actions—the fact that they were different types of actions requiring synchronisation—meant they were unlikely to have been distributed successfully if each chose his or her action independently of the others. So each made inferences about the others' reasons, beliefs, and decisions and used those inferences to help form his or her own decisions.

No individual performed the rescue. To say that any one of them performed the rescue would be to do disservice to the others' agency, like the six-year-old who does a disservice to his father's agency when he insists that *he* baked the cake. The rescue was the result of the autonomously formed decisions of separate individuals, each of whom responded to the others as features of the environment. If all but one individual's actions were fixed—say, by laws of nature—then we could perhaps say the remaining individual “performed the rescue.” But the others' actions were not so fixed. They were chosen. Each chose to respond to all the others in an on-going way. Each performed only his or her individual responsive actions—though each did so with the hope and intention that the swimmer would cease drowning as a result.<sup>36</sup> No individual performed the rescue.

Assuming that no individual performed the rescue, there are two interpretations of the beachgoers' actions. The first holds that only agents can act and that each agent has a decision-making procedure. “Only agents can act” implies not just “only an agent can do any component of any action,” but also “only an agent can do any action as a whole.” That is, if rescuing the swimmer was an action, then only an agent could have done it. The only decision-making procedures at the beach were those of each individual, so the group was not an agent.

Thus, on this first view, the group did not perform the action of rescuing the swimmer. Neither did any individual perform that action. The rescue, then, was not an action. However, “where there is an action which some individuals perform, those individuals make up an agent” is consistent with “where there is an outcome which results from some individuals' actions, those individuals need not make up an agent.” On this first

---

<sup>36</sup> I intentionally say “cease drowning” rather than “be rescued” or “be prevented from drowning.” The latter locutions imply the existence of an agent of the rescue or prevention. But the existence of such an agent is precisely what is under question.

view, “the swimmer is no longer drowning” was an *outcome* that resulted from the conjunction of five people’s *actions*. While we might colloquially say “the beachgoers rescued the swimmer,” the non-drowning outcome was rather a happy consequence of a (intentional) conjunction of five people’s (mutually responsive) actions. Each person *intended* to have some role in producing this outcome, each *hoped* that this outcome would result, but none *controlled* the outcome.

The second interpretation says that the responsive beachgoers, as a group, rescued the swimmer. The rescue is potentially an example of: Michael Bratman’s (1992, 328) “shared cooperative activity,” in which “each participating agent [knows] that the other is attempting to be similarly responsive”; Christopher Kutz’s (2000) acting together through “participatory intentions” to do one’s part; or Raimo Tuomela’s (2006, 38) joint intentional action, which “amounts to the group members’ jointly intending X and jointly acting on the basis of their joint intention.”<sup>37</sup> None of these entail—though none are incompatible with—a group-level decision-making procedure. If the beachgoers did not constitute an agent yet the rescue was an action, then this view denies that only agents can act. This is motivated by the intuition that the swimmer was intentionally rescued, and so must have been rescued *by* something.

We must weigh up this intuition against the intuition that only agents can act. As I see it, that the first view can account for the rescue—it was an outcome aimed at by several individuals, but not an action—weakens the pull towards denying that only agents can act. Moreover, one who holds the second view is committed to the existence of innumerable acts, performed by only minimally responsive groups of individuals just so long as each individual intends that some distant outcome will result and hopes she has some part in causing it. Many of these outcomes look far too distant from the agents to count as their actions.

However, someone who believes the beachgoers rescued the swimmer might claim that they constituted an agent. If they constituted an agent, then the rescue could have been an action. This would require that a group can be an agent despite lacking a group-level decision-making procedure. So let’s consider what reasons there might be for why a group-level decision-making procedure—which these beachgoers lack—is necessary for a group to have agency.

---

<sup>37</sup> Tuomela (2006) takes his joint intentional action to be more than just individuals’ sharing an aim and taking each other as features of the environment in acting to achieve that aim. As will become clear in §3.2.4, however, his “joint action” is weaker than my “collective agency.”

### 3.2.4 *Collectives*

Imagine a different group of beachgoers, identical to the first except that each beachgoer is wholly unknowledgeable about rescuing swimmers—except Laura. All see the drowning, but only Laura knows what any of the required individual actions are. In fact, she knows what *all* the required actions are. Laura asks if any of the others know what to do, and receives puzzled looks in response. So she starts instructing one to drag the boat, one to pull the starter cord, and so on. At each instruction, Laura checks that the relevant beachgoer is willing and able to follow the instruction. Each commits (if only tacitly) to following her instructions and supposes the others have too. Laura’s instructions divide the necessary actions among the sunbathers. Each performs the action that Laura instructs him or her to perform, because Laura has instructed it. The swimmer ceases drowning.<sup>38</sup>

Here, Laura has served as the group’s decision-making procedure. This example shows that creating a collective can be extremely easy. A collective has a group-level decision-making procedure—a process that takes in reasons and produces aims and instructions—that is *operationally distinct* from the procedures held respectively by its members. It is operationally distinct in the following way: its reasons and beliefs (inputs) may not be identical to the reasons and beliefs of any members; its method for processing inputs is different from the method of any one member when deciding for herself; and the decisions it produces are not the straightforward conjunction of individuals’ decisions. But the members all have some influence, even if only in the way the unknowledgeable beachgoers do—i.e., that the leader procedurally asks if they have any inputs to make. If one of the other beachgoers had tried to make decisions with Laura, then the group’s decision-making procedure would have taken on a much more complex form.<sup>39</sup>

A decision-making *procedure* has not been established just when a group reaches a substantive decision. When individuals merely agree (make a decision) on individual actions or shared goals, this may be mutual responsiveness rather than the formation of a collective. The key feature of the decision-making *procedure* is that members expect to be

---

<sup>38</sup> See fn. 36.

<sup>39</sup> As another example of operational distinctness, consider a three-person majoritarian democracy that is deciding whether  $(P \& (Q \& R))$ . If A believes  $(P \& (Q \& \sim R))$ , B believes  $(P \& (\sim Q \& R))$ , and C believes  $(\sim P \& (Q \& R))$ , then, if the group votes on each of P, Q, and R in turn, the group will believe that  $(P \& (Q \& R))$  even though none of the members believe this (List and Pettit 2011, 43–7). Its method for reaching the decision (i.e., voting) is distinct from that of each member, and the decision itself is not just the conjunction of the members’ decisions. More generally, Kenneth Arrow (1963) famously showed that, under certain conditions, *any* decision-making procedure for aggregating individual judgments into group judgments (when there are more than two possible group judgments) must treat some individuals or issues as more important than others, or must let the group’s view on an issue be determined by the group’s view on other issues. The point is that some groups can—indeed, must—make decisions that are not merely the sum of members’ decisions.

applicable to at least one *future* decision.<sup>40</sup> Examples of such procedures include: consensus among group members, majority voting, tossing a coin, one member being an authority figure, instituting sub-groups with designated fields over which the sub-group members must reach a decision by some specified procedure, and so on.

For my purposes, a “collective” is a group of individuals that have a shared decision-making procedure. A group has such a procedure if:

- (i) there are members, i.e., agents that: (a) are given roles by the procedure in (ii); and that (b) according to the procedure in (ii), are permitted by their role to have influence on the procedure’s results; and
- (ii) there is a procedure that is operationally distinct from the analogous procedure of each member, which is able to, and which each member expects to, generate decisions on the following:
  - (a) which action/s the collective will aim to perform, and
  - (b) individual roles that are jointly sufficient to perform those action/s, and
  - (c) a division of the roles in (b) between members; and
- (iii) each member commits to abide by the procedure’s decisions; and
- (iv) each member is able to receive sufficient information about the procedure’s decisions for them to abide by it.

A few clarifications. First, in (iii), being committed requires that members abide by the procedure’s decisions *because* they result from a procedure to which they and sufficient others have committed (and they know this). So members might be mistaken about whether they are in a collective if they incorrectly assume sufficient others have committed. Second, being committed does not require that one *never fails* to abide by the procedure’s decisions—the commitment might be outweighed or undermined by other reasons individuals have, though the fact of commitment suggests that the outweighing or undermining reasons would have to be all the more weighty. Third, one might “abide” without acting, for example if a collective aim is to let  $\varphi$  happen, and  $\varphi$  will happen if and only if no member acts. Fourth, not all members need know what all the collective’s decisions are. Finally, to meet (iv), intra-collective communication might be only unilateral.

Conditions (i)–(iv) are sufficient for group agency. But they may not each be strictly necessary. We could imagine, for example, a large group in which some members don’t

---

<sup>40</sup> Expectation is all that is necessary because the procedure might change before a future decision is made, but this does not mean that there was no procedure in place.

have all the expectations in (ii), or in which one member abides reliably but is not committed. A full defence of this precise account of collective agency would be tangential to the overall purposes of this thesis. But it suffices as a sketch of the requirements a group must meet in order to have agency.<sup>41</sup> Though it is operationally distinct from its members' procedures, the collective's decision-making procedure is constituted entirely by procedures that individuals use to interact with one another. A collective is nothing more than individuals arranged in a certain way—namely, arranged in a way that adheres to the procedure by which they have all committed to abide.<sup>42</sup>

The decisions produced under such arrangements are made by the collective, rather than the individuals. But we should not take this to mean that these decisions are not reducible to facts about the individuals. As an analogy, consider an open box. It is nothing more than atoms arranged in a particular way—namely, arranged in an open-box shape. Yet the open box can sit in water such that the box does not contain any water, even though none of its atoms can do this. And this property of “not containing water” doesn't derive *straightforwardly* from adding together properties that its atoms have individually (unlike, say, the box's mass). Rather, the property of “not containing water” can only be reduced to the atoms' properties by going into detail about the way the atoms are situated *relative to one another*. The same goes for collectives' decisions: the way individuals' properties

---

<sup>41</sup> My conditions are similar to some extant accounts, though nothing much rides on one's endorsement of my accounts over these other ones. First, Philip Pettit and David Schweikard's (2006, 33) conditions for “group agents” are:

[f]irst, the members act jointly to set up [which might minimally be ‘each commits to’] certain common goals [a minimal goal might be ‘that the group instructs members’] and to set up a procedure for identifying further goals on later occasions. Second, the members act jointly to set up a body of judgments for rationally guiding action in support of those goals, and a procedure for rationally developing those judgments further as occasion demands. And third, they act jointly to identify those who shall act on any occasion in pursuit of the goals...

Second, Peter French's (1979, 212) “corporate persons” must meet conditions that are, in a way, more permissive: “(I) an organizational or responsibility flow chart that delineates stations and levels within the corporate power structure and (2) corporate decision recognition rule(s) (usually embedded in something called ‘corporation policy’).” [Mix of Roman and Arabic numerals in the original text.] But French suggests that the organizational charter and corporate internal decision-making structure must be written down or otherwise formalized. In this way, his account is more restrictive than mine, which requires no such formalization.

Third, List and Pettit (2011, 35) say that the members of a group agent are those who do at least one of two things: (1) “authorise” the group to act for them; (2) play an “active” role in the pursuit of the group's ends.

<sup>42</sup> List and Pettit (2011, ch. 3) gives a survey of some of the different forms collective decision-making procedures (or, as they call them, “aggregation functions of individual attitudes”) can take. These include having group decisions track just one member's decisions, having them track the median of all members' decisions, having them constrained by vetoes and anti-vetoes, having the set of group decisions at the lowest possible “distance” from the “furthest away” (most different) set of individual decisions among members; having sub-groups “specialise” in particular decisions and letting those groups dictate those decisions, or having the group's later decisions depend on its previous decisions. List and Pettit favour the latter three types.

(reasons, decisions, commitments, beliefs, etc.) sit *relative to one another* are crucial for reducing the collective's decision to properties about the individuals.

This does not necessarily make such reduction impossible—although it does make it very difficult; perhaps too difficult for humans to give fully.<sup>43</sup> As List and Pettit (2011, 77) point out, the decision-making procedure might be non-explicit (such that individuals revise earlier decisions of the collective on the basis of new ones); collectives' decisions might be multiply realisable by individual decisions (so that a determinate reduction to individual decisions is not possible); and a collective's decision on some issue might be determined not by individuals' decisions on *that* issue, but by their decisions on numerous earlier issues, where those earlier issues (combined with the collective's decision-making procedure) entail the collective's decision on the later issue.

If collectives are really just individuals arranged in a certain way, then why think collectives are agents while sets of responsive individuals are not? Simply, because the latter do not have a shared decision-making procedure. I said earlier that agents form decisions by processing reasons and beliefs. The only groups that can form decisions in this way in their own right—in a way that is distinct from their members—are collectives, i.e., groups with decision-making procedures. Collectives do this in a way that is distinct from their members because their procedures have distinct inputs and processing mechanisms, even though the procedure depends entirely on facts about the members (facts relating to the members' commitments, expectations, inputs to the collective procedure, etc.). If my gloss on “agent” is correct, then collectives are the only group agents.<sup>44</sup>

---

<sup>43</sup> List and Pettit (2011, 4–5) also argue that group agents have “reality” and “significance” while maintaining that the “agency of group agents depends wholly on the organization and behaviour of individual members.” Yet their view’s relationship to mine is not entirely clear: they maintain that “talk of group agents cannot be dispensed with in favour of talk about individual agents,” and it is unclear whether they believe it cannot be dispensed with because such dispensing would be *difficult* (as I claim) or because it would actually miss out some facts (which I deny). Their claim that the group is “not *readily* reducible” (emphasis added) to individuals suggests that such reduction would merely be difficult. Yet they claim that viewing group agents as intentional subjects (which seems to require viewing them non-reductively) allows us to “understand them in a new way” (2011, 11–12) and that by refusing to conceptualise the group’s doings at an individual level we “gain a perspective that opens up new possibilities of interacting within the social world” (2011, 76). The latter suggests that there is something substantially lost in any attempt at a reduction. More importantly, they later claim that *both* the collective *and* its members can be held retrospectively responsible for some action of the group’s: “we should hold the various enactors [members] responsible, if circumstances allow, for any harm their voluntary acts and omissions produce. But we should also hold the corporate entity responsible for the harm that it arranges to have done, given the decisions it licenses and the procedures by which it channels those decisions.” But if the reduction of groups’ actions to individuals’ actions is possible (just not “readily” so), then the collective’s responsibility should be reducible to individuals’ responsibility. If this is right, then the collective does not “arrange,” “license,” or “channel” decisions or actions in a way that is not reducible to facts about individuals. There is then no need to hold the collective itself responsible, as List and Pettit would.

<sup>44</sup> Some will deny that decision-making procedures are sufficient for agency; or that any groups can have distinct decision-making procedures. From here, my arguments are directed at those who agree that collectives have minimal agency. Conversely, one might hold a more permissive account than mine of the conditions a group must meet to act, bear duties, or be an agent. For example, one might maintain that

Collectives' acts are constituted entirely by (conjunctions of) individuals' acts, and collectives' reasons and beliefs (the procedure's inputs) must first be produced by individuals recognising reasons, or forming beliefs, and then inputting those according to the collective's procedure. While I have focused on *decision*-making procedures, collectives will also themselves need reason- and belief-forming procedures, the outputs of which will be inputs in the decision-making procedure. At some point, this regression of collective procedures will stop at a procedure which has inputs entirely produced by individuals' own procedures. But collectives' decision-making procedures can process reasons and beliefs, and produce instructions for members, in ways that enable members to act not just on their own decisions, but as members of a collective. That is, they enable members to act in a manner sanctioned by a decision-making procedure that is different from his or any other member's procedure—in a manner sanctioned by the collective agent.

I will take it that an individual *performs his role* in a collective act when and only when conditions (i)–(iv) above are met and:

- (v) he is assigned a role, R, by the procedure in (ii). R is a member of a collection of roles that are jointly sufficient for performing a collective act which the procedure in (ii) has decided to perform.
- (vi) he acts within R with a view to the collective act's being performed.

A collective act is the carrying out of one of the decisions in (ii)(a), through members acting within the role division in (ii)(b) and (ii)(c). We can attribute an act  $\varphi$  to the collective for three reasons. First, performing  $\varphi$  was a distinctively collective aim—the decision to aim at it was produced by the collective's procedure. Second, the sufficient means for performing  $\varphi$  was determined and distributed to its members by the collective itself. Third, the members played the roles they did towards  $\varphi$  because the collective so instructed them.

With this conceptualisation of group agency in place, we can turn to the issue of collectives' duties.

---

groups of responsive individuals can act or are agents. Such a person will still find my account of collective duties of interest, since he can view it as an account of the distribution and implications of the duties of whichever groups he thinks can have duties.

### 3.3 Collectives' Duties and Coordination Duties

#### 3.3.1 Collectives' Capacities and Dependency Duties

I began this chapter by noting that agents can only have duties to take measures available to them—they cannot have duties to take measures available to others but not to them. Nonetheless, human moral agents are able to recognise the moral reasons that apply to agents other than themselves. If a collective's members are human moral agents, and so can recognise moral reasons, then it is natural to assume they can design a collective procedure into which they can put the moral reasons the *collective* should respond to. And it is natural to assume they can design the procedure so that the procedure processes the moral reasons in the way morality demands of the collective, such that the collective takes its *own* measure in response to those reasons.

If all of this is possible, then a collective can take measures in a way that responds to those reasons, given that members designed the procedure in the appropriate way and put its moral reasons into the procedure at the appropriate point. All of this is to say that collectives can bear and discharge duties. From here on, I will focus on collectives that can form decisions, and distribute roles, in response to their duties. This is probably not true of all collectives (probably not all collectives are set up such that they can process moral reasons), but these are the collectives I am interested in.

In Chapter Two, I said that the following is a necessary condition for an agent to have a dependency duty to take some measure: the likelihood that the interest will be fulfilled if the measure is taken is *proportionate* to the importance of the interest, where more important interests are proportionate to a wider range of likelihoods (with a lower bound that is lower). If an agent's most efficacious measure for fulfilling an interest meets this condition, let us say he (she, it) is *sufficiently capable* of fulfilling the interest. In assessing whether collectives are sufficiently capable of fulfilling some interest (so as to be eligible for a dependency duty), we must consider the collective's most efficacious measure.

Generally speaking, the measures that collectives can take to fulfil interests are the following: first, deciding to take measures to fulfil the interest; second, distributing roles to members jointly sufficient for fulfilling the interest (at least, sufficient in many likely futures); and third, enough members acting within their respective roles with a view to fulfilling the interest.<sup>45</sup> If these three measures would be insufficiently likely to fulfil the

---

<sup>45</sup> I say “within their roles”—i.e., “acting consistently with their roles”—rather than “performing their roles” because I mean to be permissive about what individuals might do (and be required to do) in pursuit of the collective’s fulfilment of an interest. Acting consistently with one’s role might require more than merely performing the role one was explicitly given (such as “drag the boat”). For example, it might require taking up

interest if the measures were taken, then the collective is insufficiently capable to have a duty to take the measures. That is: even if a collective meets (i)-(vi), it cannot have a dependency duty to take measures to fulfil an important interest if it is insufficiently likely that the conjunction of these three measures will fulfil the interest. (Of course, to have a dependency duty it must also meet the other conditions of the Dependency Principle, regarding the expected realisation of positive value, the expected realisation of no less value than other agents, and so on.)

When we are determining whether an agent is sufficiently capable of fulfilling interests to have a dependency duty, we do not let her unwillingness to take the measures impede the possibility of her having a duty to take them. But one might query whether we should make this assumption about the third collective measure—the measure of enough members acting consistently with their respective roles with a view to fulfilling the interest. After all, members often won't do this. And their unwillingness is hardly the collective's fault. So perhaps we should not assume that the members will perform their roles, when figuring out whether a collective is sufficiently capable. Perhaps we should rather treat members as objects of the collective's environment, to be worked around. When we were defining individuals' capacities, we were not allowed to simply assume that others would do their duty—rather, we were to assume that the agent in question will do their duty, but be realistic about the likelihood that any other agent would do theirs. Surely the same should hold for collectives: the members of a collective are *different agents* from the collective itself, so we cannot assume they will comply with morality's demands when identifying the collective's duties.

However, it is unclear what it would even mean to assume a collective will take measures without assuming that any of its members will take measures: the collective's measures are constituted by those of individuals who are arranged in a certain way. The assumption that the collective is willing entails the assumption that (enough of) the members are willing. Moreover, since condition (ii) for collectives says that the individuals have committed to the procedure, the assumption that they will try is natural when we are

---

the slack for other members' failures, taking the initiative in reaction to unpredicted environmental obstacles, and so on. When a collective has a dependency duty, it must distribute roles that would be *sufficient* for interest fulfilment in *many* ways the world is *likely* to be, *assuming* that all members will do as they are told. But if the world or other members do not cooperate, then “acting within one's role” to fulfil an interest might require more than simply doing what one was told. That said, acting *inconsistently* with one's role in order to fulfil the relevant interest would not help the collective to discharge its duty. For example, a humanitarian worker might consider bribing local warlords to perform her role of “getting food aid to refugees.” But suppose the collective has prohibited her from offering bribes. If she can bribe the warlords to get the food to the refugees, then *she* has measures open to her for fulfilling the refugees' important interests—but the *collective* does not have such measures open to it. Her taking such measures would not contribute to the *collective's* discharging *its* duty.

thinking about whether the collective's potential measures will be sufficiently likely to succeed to generate a duty to take them.

Despite this tight connection between collectives and their members, collectives often have dependency duties when individuals do not. This is because the likelihood of their measures succeeding is often higher than the likelihood of any given individual's measures succeeding—or even of a number of individuals' aggregated measures succeeding. To see this, simply recall the unknowledgeable beachgoers. Without the collective decision-making procedure—without someone making decisions and distributing instructions *for the group*—success would have been extremely unlikely. This would be true even if each individual were responsive to the others. If the individuals had tried to be responsive, the swimmer would not have been saved.

Thus the aggregated capacities of the set of responsive but unknowledgeable beachgoers did not simply “add up” to the capacity of the Laura-led unknowledgeable rescue team: the responsive individuals, taken in aggregate, did not have such a high likelihood of rescuing the swimmer if they took their most efficacious measures. Assuming the sufficient capacity for having a duty to rescue the swimmer is anything higher than “extremely low,” it is only with the collective decision-making procedure that measures were sufficiently likely to succeed for there to be a *duty* to rescue the swimmer. This shows that collectives’ capacities are not reducible to the capacities their members have when there is no collective decision-making procedure, that is, that the members would have independently of the procedure.

However, collectives’ capacities *are* almost entirely reducible to the sum of the capacities members have when there *is* such a procedure. It is tempting to picture the procedure as a kind of independent machine, which churns out decisions and roles of its own accord. It has, as it were, a mind of its own, and members have merely committed to follow the instructions it autonomously produces (Pettit (2010) evokes precisely this image). But collectives are not really like this. Their procedures rely on individual members for the procedure to be implemented. The machine needs someone to turn the crank (not to mention stock the machine with materials, e.g., beliefs, desires, moral reasons), and only members can do this. The collective structure is a way for individuals to organise themselves; it is not some external force that acts upon them or independently of them. Collective structures organise individuals so that they can do more than they could *without* the organisation, but the collective can do little more than the sum of what individuals can do *within* the organisation.

For example, when we say that the *collective* distributes roles, we mean that *individuals* use the collective procedure, to which they and a number of others have committed, to

distribute roles. The individuals' capacities to distribute roles do not exist without the collective, but they are capacities of individuals. Similarly when we say that the *collective* decides to  $\varphi$ , we mean that individuals make their respective contributions to the procedure, and then operate the procedure to calculate their shared decision.

Collectives' capacities are thus almost fully reducible to the sum of individuals' capacities in collective contexts—but not quite. In collectives—unlike in non-collective groups—there is a locus of agency (the group's decision-making procedure) that can (as long as its handle is cranked) *produce multilateralism* among members. The collective's capacity to produce multilateralism in the pursuit of collective goals is analogous to my capacity to produce multilateralism in the parts of myself to run 10km. My legs are capable of moving the right away, my lungs are capable of breathing the right way, and so on. My capacity to run 10km is, in some sense, nothing more than these capacities.

But neither my legs, nor my lungs, nor any other part of me has the capacity reliably to produce multilateralism among the legs, lungs, etc. in order to actually produce a 10km run. Only I can produce the multilateralism. The analogous capacity of collectives to generate *multilateral* role performance among members is the “remainder” of the collective's capacity—this is the bit that is not quite reducible to the members' capacities, since it is the collective machine itself (rather than members acting within that machine, or turning the machine's handle) that produces the multilateralism. Such multilateralism is essential to, for example, the Laura-led rescue team's capacity to rescue the swimmer.

The capacity to produce multilateralism is a capacity of the collective framework, as those frameworks were schematised in conditions (i)–(v). The collective framework is reducible to members' reasons, decisions, commitments, beliefs, and so on. Thus although the capacity to produce multilateralism is not reducible to the sum of members' *capacities*, it is reducible to certain *facts* about the members. Specifically, it is reducible to the facts about them that constitute the collective framework: their beliefs, commitments, expectations, and so on, as outlined in (i)–(iv). (I said that (i)–(iv) may not be precisely correct. But whatever the precisely correct conditions are, I assume that they will refer only to certain facts about individuals, such as their beliefs, commitments, and so on.)

This is important: if collectives have duties based on being best-placed to fulfil an important interest (i.e., if they have dependency duties), *and* if collectives' being best-placed is wholly re-describable in terms of facts about individuals (whether individuals' capacities, or individuals' commitments, beliefs, and so on, which make up the collective framework), then some of collectives' duties (specifically, their dependency duties) are explicable in

terms of facts about individuals. Collectives' dependency duties are thus explicable via facts about those collectives' members.<sup>46</sup>

Yet we should not do away with talk of collectives' capacities and dependency duties in favour of wholly individualistic talk. This is because of the descriptive power collective talk gives us. For example, assume the beach rescue team's capacity to rescue the swimmer is nothing more than Laura's capacity to choose and distribute tasks, and the others' capacities to perform those tasks once given them. The sum of these capacities is the capacity to rescue the swimmer. But because no individual has *that* capacity on their own, talk of the group agent's "capacity to rescue the swimmer" is helpful shorthand for talk of the individuals' capacities to do various bits of that group agent's action. It would be extremely complicated to reduce the collective capacity not only to individuals' capacities, but also to their commitments, beliefs, and so on which constitute the multilateralism-producing collective framework. (It is not nearly so complicated to do this reduction for the outcomes of responsive individuals, which is why I suggest we talk only about individuals in those cases.)

Given that we are concerned with dependence-based duties, it makes sense to ask what features of a collective tend to make it more capable, and more likely to be best-placed to fulfil an interest. Multiple features come into play here. These include: the maximum sanctions it can impose on members; the maximum percentage of their resources it can commandeer; the number of possible futures in which the decision-making procedure will apply (that is, be able to make decisions); and the likelihood that causes external to the group, such as other agents or natural events, will inhibit the decision-making procedure's ability to achieve its aims. This list is not exhaustive, but should give some indication of the things that a collective (or an aggregate of individuals that are setting up a group—on which more below) should attend to, if it is to be as capable as possible of performing some action.

If I am right that collectives have moral agency, then collectives can have duties. If I am right that collectives can be better-placed to do things than any individual is alone, then collectives can have dependency duties. This increases the scope of the Dependency Principle, by increasing the number of duties it produces. However, there remains a significant problem for ascribing dependency duties to groups.

---

<sup>46</sup> While this works for dependency, it may not work for other duty bases: perhaps collectives can cause harm, benefit from injustice, make promises or sign contracts, and so on in ways that are not explainable wholly in terms of facts about individuals. If that can happen, then those other kinds of duties will not be explainable in terms of facts about individuals.

### 3.3.2 Coordination Duties: The Problem in General<sup>47</sup>

The above analysis of collectives' capacities will have bearing on all sorts of collective duties, with all sorts of bases. Yet there is a problem for group duties that seems most pertinent to dependence-based duties, as against duties of contract, harm, association, and so on. This is because the moral imperatives for non-dependence-based duties tend to arise out of some property that the duty-bearing agent *already* has. Often, these properties are had in virtue of a past action—causing harm, signing a contract, developing a relationship, or similar. There are exceptions: duties to not-harm, for example, do not require some past action on the part of the duty bearer, and might be held even if the agent is unable to do harm. But this duty has no *practical upshot* unless the capacity to harm already exists. The standard range of action-demanding non-dependence-based duties are grounded in facts that are *already* true of some agent: they have backward-looking, rather than forward-looking, justifications.<sup>48</sup>

Dependency duties are grounded in an agent's being best-placed to take measures that would, if taken, fulfil an important interest. This is usually a property the agent already has. Yet in some cases the very measures that seem, intuitively, to generate the duty cannot be taken by an already-existing agent. In these cases, the capacity to take interest-fulfilling measures does not exist at all. Rather, a number of already-existing agents each have the capacity to take measures to be responsive to one another with a view to fulfilling the interest. But none of them has the capacity to fulfil the interest.

We are tempted in such cases to say the *aggregate* of individuals has a capacity, and therefore has a duty: “the hikers have a duty to lift the fallen tree off the child,” “the pedestrians have a duty to stop the murderer,” “the beachgoers have a duty to save the drowning person.” (For simplicity, in each case suppose an important interest can be fulfilled *only if* those particular individuals work together: these individuals are *uniquely* well-enough placed to bear a duty. So, the duties derive from being uniquely well-enough placed, which is one way of being best-placed (i.e. being uniquely well-enough placed is one way of meeting the Dependency Principle).)

These cases are puzzling. The person whose interest is unfulfilled is in a very real sense *dependent on* each member of the aggregate, and each member of the aggregate intuitively has a duty to contribute to the interest's fulfilment. These duties seem to derive

---

<sup>47</sup> The remainder of §3.3 expands on the arguments in Collins 2013.

<sup>48</sup> Though there is sometimes an issue with picking out the *agent* who bears the duty-generating property, for example when a non-collective angry mob seems intuitively to have caused harm and to owe compensation to victims, but we cannot identify precisely which agent bears the “caused harm” property. Yet this problem with identifying the bearer/s of non-dependence-based duties is distinct from the point that these duties look to already-instantiated properties, rather than possible-but-not-yet-instantiated ones.

from their being best-placed (because uniquely well-enough placed) to do so, in some sense. Yet if the hikers, pedestrians, and beachgoers are not collectives, then they have no agency as groups. There is no *agent* that is capable of fulfilling the interest. How, then, are we to understand this dependence relation, and the duties that derive from it?

Ample attention has been paid in recent years to the conditions under which a collective agent exists (List and Pettit 2011; Pettit 2007, 2010; Pettit and Schweikard 2006) and the conditions under which it can be said that people are “acting together” (Bratman 1992; Gilbert 2001, 2006; Sudgen 2000; Tuomela 2006). Yet surprisingly few have taken up the issue of when and why individuals might severally bear *duties* to act together or to take steps towards *creating* a collective agent.<sup>49</sup> Accounting for these duties is a very important problem for the Dependency Principle, as there are very many situations in which dependence intuitively generates duties, but where there is no already-existent agent that is sufficiently well-placed to bear a duty. These include the situations of the hikers, pedestrians, and beachgoers just mentioned, but also (as we will see in Part III) much more large-scale cases, such as the situations of aggregates of states who together—but only together—have the capacity to protect the world population from mass atrocities. I will address such large-scale cases later. But it will be helpful to begin by analysing some simpler ones.

### **3.3.3 The Problem for Duties to be Responsive**

Go back to the beachgoers. Suppose Ben, Jon, Julie, and Stan are oblivious to the drowning and just happen to be taking the boat to where the swimmer is. Laura is strong enough to pull the swimmer to safety alone. She can respond by getting in the boat and doing the pulling. Imagine that the Dependency Principle’s antecedent holds: the swimmer’s important interest in not drowning today is unfulfilled; Laura is capable of taking measures (getting in the boat and pulling) that have a sufficiently high likelihood of saving him; these measures would realise positive expected value regarding Laura and the swimmer; doing the same for others would realise positive expected value regarding the

---

<sup>49</sup> One exception is Virginia Held (1970, 480), who comments in passing that

... it may well follow that in some cases all the individuals in a random collection are responsible for not acting to transform the collection into an organized group, even though none of these individuals is responsible for not taking the action that ought to have been taken by an organized group in these circumstances.

However, she doesn’t explain any further and the rest of her article defends ad-hoc groups as duty-bearers, which I have just denied. Holly Lawford-Smith (2012, 6) suggests that “four strangers who happen to be passing in the stairwell ... [where a piano has been dropped on a child] might be blameworthy if they do not attempt to *form* the kind of collective that could lift the piano,” but worries that “this is plausibly itself a collective action” and that “non-existent agents cannot have obligations.” Later, I address these concerns.

relevant parties; and her measures would realise no less expected value regarding agent and swimmer than any other agent's most efficacious measure.

She then has a dependency duty to jump in the boat and pull the swimmer in. Her duty is based on the fact that she alone, by treating others as environmental objects and acting responsively to them, is best-placed to rescue the swimmer—which entails that others' actions are sufficiently fixed. If she succeeds, the action is not a group action, but an individual action. Her duty is a dependency duty.

Of course, her responsiveness is likely to fail if she makes an incorrect guess about how others will behave, or what she should do in response. If such slip-ups are quite likely, then Laura will have a low likelihood of rescuing the swimmer if she takes her measures. If this likelihood is low, then the swimmer's interest—his interest in life—will have to be very important in order for Laura to be sufficiently capable to have a dependency duty to take those measures. But this is perfectly possible, since life is plausibly the *most* important interest. Thus it is plausible that the Dependency Principle applies to Laura and to Laura alone.

In this way, cases that require just one agent to be unilaterally responsive to other agents are covered by the Dependency Principle, since the other agents are treated just like predictable objects in the environment. This kind of case is not problematic for the Dependency Principle.

Things get problematic if assistance necessitates responsiveness that is more than minimal—such as mutual responsiveness between several individuals with overlapping, commonly known intentions. Plausibly, if such responsiveness is sufficiently likely to result in the non-drowning outcome, and nothing else is sufficiently likely, then each individual beachgoer has a duty to be responsive in this way. Such duties seem to exist in the first beach drowning case, where each beachgoer is knowledgeable. Moreover, these duties are intuitively grounded in the swimmer's dependence on the beachgoers.

Yet the Dependency Principle cannot generate these duties, as it is false of each beachgoer that he or she is sufficiently capable of fulfilling the swimmer's interest in being rescued. (That is, no beachgoer meets condition (2) of the Dependency Principle.) Each beachgoer is capable of “through responsiveness, making a contribution to the non-drowning outcome's being realised.” The contributions in aggregate have a high likelihood of producing the non-drowning outcome, but no one contribution has a sufficiently high likelihood of producing that outcome.

We might be tempted to say that *the aggregate* has sufficient capacity to rescue the swimmer. But if this is a capacity of the aggregate, then whose duty is it to intentionally take the relevant measures? It cannot be the aggregate's duty, as it has no agency so it

cannot intentionally do anything, including taking interest-fulfilling measures. And it cannot be the duty of any of the individuals that make up the aggregate, as they cannot take measures that are not theirs. While it might be true, in some sense, that the aggregate has the capacity, *this* capacity cannot be the basis of a duty, because it is not the capacity of any agent that can *exercise* the capacity, i.e., that can take the relevant measures.

We might rather say that each individual has the capacity to *make a contribution* to the non-drowning outcome, and that each of them incurs a dependency duty because that contribution is itself something in which the swimmer has an interest. Recall the Dependency Principle requires that (1) B has an important final interest that is unfulfilled; and (2) If A intentionally takes measure M, where M is the most efficacious measures open to A to fulfil the interest in (1); then, even if other agents do not coordinate with A, then FI is likely to be fulfilled with a likelihood that is *proportionate* to FI's importance, where more important interests can trigger duties on the basis of lower likelihoods. Does B really have an important interest in *each* individual's contribution, taken without regard to the others' contributions? B has a vital interest in *all* the beachgoers doing their bit; not in *each* of them doing their bit. It is no use to the swimmer that Stan starts the boat, for example, if no one is there to jump in it. Given that each contribution is worthless without the others, it is odd to say that each, on its own, fulfils an important interest.

Another option would be to consider each beachgoer in turn. As we do so, we hold the others' responsiveness as determined. That is, when we consider knowledgeable Ben's capacity to rescue the swimmer, we hold fixed that Laura, Stan, Jon, and Julie will be responsive to each other (and to Ben) with a view to the non-drowning outcome. We treat Laura, Stan, Jon, and Julie like determined objects in the environment. This gets us the result that knowledgeable Ben has the capacity to rescue the swimmer, in the same way Laura does in the case where the others happen to be taking the boat out. We then do the same for the other four knowledgeable beachgoers (*mutatis mutandis*), with the result that each is capable of rescuing the swimmer. They then each have a dependency duty to exercise that capacity.

The problem with this is that we are not, in general, allowed to simply *assume* that the world will cooperate when we are trying to figure out whether an agent has the capacity to do something. We are trying to determine the *duties* of each beachgoer, not identify the most ideal way the world could be. We should of course account for the *probability* that the world will cooperate—this is what we did in the case where Laura had a duty to jump in

the others' boat and rescue the swimmer. But we cannot just assume that this probability is high.<sup>50</sup>

Suppose the likelihood of others' cooperation is not high. Suppose the others cannot be bothered acting responsively with a view to the non-drowning outcome, and will do so if and only if he or she has a duty to do so. It then seems bizarre that we can take as given that each *will* act responsively when determining Ben's duty—as if we can just assume that each has a duty of exactly the same type and basis as Ben's duty, when his or her performance of that duty partly grounds Ben's duty. If we can just assume the others each have the duty when determining Ben's duty, why not just assume they all have duties and be done with it? Yet it is unclear on what basis can we say this. We certainly cannot say it on the basis of the Dependency Principle.

### **3.3.4 The Problem for Duties to Transform Collectives**

Similar problems arise in differently structured cases. Sometimes, a collective lacks both the short-term and the long-term capacity to take some measure. If it lacks the capacity in both these senses, then the *collective* is off the hook. But the associated individuals, intuitively, certainly should not be let off the hook.

For example, maybe the beach rescue team has a diachronic dependency duty to acquire ten more members, in order to enable it to perform future rescues. Yet just like individuals, collectives must be *able* to change themselves in order to have a duty to do so. For individuals, the ability to change oneself is plausibly constrained by things like (though not necessarily limited to) logical, physical, psychological, political, and social possibility. For collectives, the same sorts of constraints apply, but their interpretation is more difficult. For example, if rescue team members can use the current procedure to distribute roles to advertise for new members, then the team appears able to transform itself in this way. Yet a social constraint might exist: advertising and other forms of persuasion or coercion may be very unlikely to attract new members. Or perhaps there is no one else for miles around, so that the team is physically constrained in its ability to get more members. A version of a psychological constraint might exist if the team has a non-amendable constitution stipulating that this team will never have more than five members.

---

<sup>50</sup> This is a crucial difference between mere responsive individuals and collective agents. The collective structure allows us to *assume* multilateralism when determining the collective's duties. Of course, multilateralism will not always *actually* occur among individuals in collectives, but then we can explain it as a failure to discharge the duties they have in virtue of being members of a duty-bearing collective. I discuss this in §3.4.

Take the non-amendable constitution case. Here, the team could not gain five members while retaining its identity—so *this team* would not have the capacity to do that. For a collective to have a duty to change itself, its current decision-making procedure must be such that members can *act within their roles within* the collective to transform the collective. (I will say more about these duties below.) If the non-amendable constitution does not allow this, then members cannot have these duties qua members of this collective.

Nonetheless, the individuals might be able to change the collective from the outside, by acting *inconsistently* with their roles. Suppose that if members act inconsistently with their roles in the decision-making procedure, by overriding the constitution, then they could make the procedure such that it could distribute roles to members sufficient for adding ten members. They can tear up the old constitution and write a new one, without the new one being put in place via the authority of the old one in any sense. That is, imagine that individuals are capable of changing the team's procedures and goals by working on them from the outside, thus making the team such that it can acquire new members. Imagine further that future swimmers have an *important interest* in the team having ten more members.

Here, no agent—either individual or collective—is capable of adding ten members to the team. Yet individuals are capable of taking actions that will, in aggregate, transform an agent (the team) so that *it* is capable of transforming itself by adding ten members. If the team is transformed so that it has this capacity, then (suppose) the future swimmers' important interests in protection at the beach will generate a dependency duty for the team to add ten members. If individuals have duties to take the initial team-transforming actions, then these duties are not held in their capacity as team members, since if they could transform the team in their capacity as members, then the team would be able to transform itself. Rather, these would be duties individuals have as individual moral agents, to transform another agent so that that other agent is able to do some morally valuable thing.

In this case, at  $t_1$  before the collective is transformed, neither the collective nor the individuals have the capacity to fulfil the swimmers' important interest in protection at the beach. So the individuals' duties to transform the team are not covered by the Dependency Principle, despite the fact that the duties seem to have something to do with the swimmers' present dependence on the individuals, and the swimmers' potential dependence on the team.

Another case is one where individuals are able to cause a collective to incur a dependency duty by their *joining* the collective. This is a type of transformation-from-the-outside, akin to the constitution-changing action discussed above. Suppose that the team's constitution allows having ten more members, and that there are ten strong swimmers

nearby who would be fine additions to the team—but only if they *all* join. If the cost isn't too high, intuition tells us that the strong swimmers have duties to join the team. These duties seem based on the dependence of weak swimmers on the strong swimmers. Again, though, suppose there is nothing any of these potential new members can do for the future swimmers on their own. It would require all of them to join the team, thereby rendering the *team* sufficiently well-placed to fulfil the swimmers' important interests. These duties are not quite covered by the Dependency Principle, since no strong swimmer on his own can render the team sufficiently well-placed to bear a dependency duty.

### **3.3.5 The Problem for Duties to Form Collectives**

Consider the unknowledgeable beachgoers who, in the example, each voluntarily take individual steps towards the formation of the Laura-led rescue team, and each follow the instructions that Laura gives them, thus acting within their role in the team to see to it that the team rescues the swimmer. Even if we can solve the problems of (1) generating duties to be responsive, and of (2) generating duties to transform or join collectives, this will not be enough for (3) generating duties for the unknowledgeable beachgoers to form a collective. The non-drowning outcome is extremely unlikely to result from their being *responsive* with a view to the non-drowning outcome—it is much more likely that chaos will ensue—so intuitively they have no duties to do that. There is no extant collective in the vicinity, so they cannot have duties to *transform* (including to *join*) any extant collective. And the group of unknowledgeable beachgoers as such cannot bear a duty to rescue the swimmer, because it is not an agent.

This cannot be the full story. Someone is drowning. If the beachgoers formed a collective with Laura as leader, then that collective would incur a dependency duty to save the drowning person. In the original example, the unknowledgeable beachgoers *voluntarily* took the steps necessary to form such a collective. But we need there to be *duties* for them to do so. Surely, it goes against the spirit of the Dependency Principle that each member of an aggregate of individuals can avoid having any duties, because only a collective would incur a dependency duty and there is no collective, when those individuals can each take individual steps that would result in a dependency duty-bearing collective. Intuitively, the individuals all have duties to take individual steps towards there being a collective that can rescue the swimmer. Yet these will only be duties to take individual steps, not duties to actually form the collective, as no individual is able to form the collective on his or her own. Again—as in the cases of transforming a collective—at  $t_1$  no individual or collective has the capacity to rescue the swimmer, so none can have a dependency duty to do so.

### 3.3.6 Solving the Problem: Ruling out some Possibilities

One might think solving this problem is simple, at least for cases where individuals are capable of being responsive in a way that will, combined with others' responsiveness, *produce a collective* that can bring about the outcome (either by transforming an existent collective, or by forming a new collective). It is instructive to see why initially plausible solutions will not work. Call the important-interest-fulfilling outcome “p.” Suppose we try to build on the proposition that “A (an individual agent) has a duty to secure the means to fulfilling his duties.” If A has a duty to see to it that p, and  $\psi$  is A’s best means to p (where  $\psi$  is A’s individual collective-transforming or collective-forming responsive action), then A has a duty to  $\psi$ . This appears fine. After all, A *can* see to it that p in the long-term: A can  $\psi$  and then A can see to it that p through the collective. Because he can see to it that p in this two-step way, he cannot avoid the duty to see to it that p by pointing out that he cannot see to it that p at  $t_1$ . At  $t_1$ , he *can* bring it about that he *can* see to it that p at  $t_2$ . So, at  $t_1$ , he is eligible for a duty to see to it that p.<sup>51</sup>

However, A often cannot see to it that p at  $t_2$ , even if (at  $t_1$ ) A  $\psi$ s and others’  $\psi$ -ing is very likely. This is because the collective’s  $t_2$  p-producing actions—call them  $\varphi$ —are not performed *by A*. At  $t_2$ , A can perform only his role in the collective  $\varphi$ -ing. He cannot  $\varphi$  himself. This is perhaps false of Laura, the rescue team’s dictator. If, at  $t_1$ , it is very likely that the unknowledgeable beachgoers will take steps to form a collective, then perhaps at  $t_1$  Laura has an individual ability to rescue the swimmer: she is able to manipulate others as if they were features of the environment, so that they perform the necessary roles. Describing this as an individual ability is peculiar, because of the ongoing possibility of others’ defection. Yet suppose we grant this. Then, at  $t_1$ , Laura is eligible for a duty to rescue the swimmer by taking responsive steps to form the group—but Ben, Jon, Julie, and Stan are not. So what can we say about their duties, and more generally about the duties of any non-collective group of individuals who do not have a clear dictator?

Consider unknowledgeable Ben. He is unable, at  $t_1$ , to make himself such that he can rescue the swimmer through the collective, since he lacks sufficient control over the collective’s decision-making procedure. If Ben—as an individual—cannot bring it about that he can  $\varphi$  through the collective, then Ben cannot have a duty to  $\varphi$  through the collective. If Ben does not have a duty to  $\varphi$  through the collective, then he cannot have a duty to take responsive collectivising actions to *form* the collective. Neither can he have a

---

<sup>51</sup> In Mark Jensen’s (2009) terms, A is indirectly diachronically able to  $\varphi$ : A can  $\varphi$  later, provided he performs an enabling action first.

duty to transform the collective, in a case where a not-quite-capable collective exists, which Ben can make capable. He cannot have these duties, because the only reason he would have had them would have been because this would have been a necessary condition for his  $\varphi$ -ing (over which he would have had a duty). But if Ben cannot  $\varphi$  *through* the formed or transformed collective, then he has no duty to  $\varphi$  and thus no duty to take the means to do so. Ben is thus not going against his duties if he refrains from transforming the collective or from taking steps to form a new collective.

The same problem arises for knowledgeable Ben and his duty to act responsively to the others with a view to  $p$  (foregoing the formation or transformation of any collective). If knowledgeable Ben and all the other knowledgeable beachgoers are responsive to one another, then the non-drowning outcome will result from their responsiveness. Yet knowledgeable Ben is not capable of producing the non-drowning outcome—he does not have sufficient control over the others and the others are not reliable enough (suppose they will be responsive if and only if they have duties to do so). He thus cannot have a duty to be responsive *in order to* fulfil his duty to produce the non-drowning outcome, because he is incapable of producing that outcome and so incapable of having a duty to produce it.

Some may yet be unconvinced. One might think an agent can have a duty to *try to* produce an outcome or perform an action he has a low likelihood of producing (performing) if he tries, as long as the outcome (action) is important enough. Then each knowledgeable individual could have a duty to try to see to it that the swimmer ceases drowning, and each unknowledgeable individual could have a duty to try to  $\varphi$  through the not-yet-existent collective. A duty to try to bring about  $p$  or to  $\varphi$  plausibly amounts to a duty to perform those actions that most efficaciously increase the likelihood of  $p$ , or of one's  $\varphi$ -ing. We then straightforwardly get individual duties to be responsive with a view to  $p$  directly, or to take individual steps to form or transform a collective such that the collective can  $\varphi$ , if this is the most efficacious way for individuals to increase the likelihood of  $p$ , or to increase the likelihood of their  $\varphi$ -ing.

However, a plausible condition for an agent being able to try to  $X$  is that he believes it is possible that he will  $X$  if he tries. If this condition is correct, then only individuals satisfying it can have a duty to try to  $X$ . Of course, in the responsiveness case, each beachgoer might believe “it is possible that I will produce the non-drowning outcome if I try by being responsive, since possibly then others will also be responsive and then I will produce the non-drowning outcome.” In the collective case, they might each believe “it is possible that I will save the swimmer if I try by taking individual steps to transform or form a collective, since possibly then others will do likewise and then I will save the swimmer.” The individual’s belief that *he or she* will produce  $p$  (or will  $\varphi$ ) if others are

responsive is false (perhaps unless she is Laura), because  $p$  is an outcome of the individuals' actions in aggregate, and  $\varphi$  is an action of the transformed or newly-formed collective. Still, each beachgoer might falsely believe this. Then, each could have a duty to try to produce the non-drowning outcome or try to save the swimmer. But these duties would not be to produce that outcome or save the swimmer, only to *try* to do so. And we are still left with the problem of cases where individuals do not have false beliefs about their own capacities.

### ***3.3.7 The Coordination Principle***

We need a new principle. This will serve to supplement the Dependency Principle in cases where an important interest can be fulfilled only if—or can be fulfilled at highest expected value for agent and dependent if—a number of individuals: (i) are responsive to one another; or (ii) take individual steps towards the transformation of an existing collective; or (iii) take individual steps towards the existence of a new collective. In case (i), the new principle will serve to replace the Dependency Principle entirely, as the fulfilment of its duties will get us straight to the fulfilment of important interests. In cases (ii) and (iii), the new principle will act as a “trigger” for the Dependency Principle: it will generate duties to create an agent that meets the Dependency Principle. I propose the following Coordination Principle:

If:

- (1) “ $p$ ” is a non-actual state-of-affairs in which the important interest of  $X$  is fulfilled; and
- (2) At  $t_1$ , either: no (collective or individual) agent has a duty to see to it that  $p$ ; or: any such agent defaults, such that  $p$  will not occur; and
- (3) Either:
  - (3a) if, at  $t_1$ ,  $A, \dots, N$  each took responsive steps with a view to  $p$ , then it is sufficiently likely<sup>52</sup> that  $p$  would occur;

or:

- (3b) if, at  $t_1$ ,  $A, \dots, N$  each took responsive steps towards there being a collective-that-can-produce- $p$  (either by transforming a current collective or taking steps towards the existence of a new collective), then it is sufficiently

likely<sup>52</sup> that, at  $t_2$ , that collective would incur and discharge a dependency duty to  $p$ ; and

- (4) At  $t_1$ ,  $A, \dots, N$  are each able to take the responsive steps referred to in either:
  - (4a) 3a, or:
  - (4b) 3bin a way that realises positive iterative expected value regarding themselves and  $X$ ; and
- (5) No other set of individuals  $\{a, \dots, n\}$  is such that each member of  $\{a, \dots, n\}$  would (if they had a duty to do so) take steps towards  $p$  or towards there being a collective that would incur a dependency duty to  $p$ , where  $\{a, \dots, n\}$  will produce higher expected value (regarding the agents and  $X$ ) through its steps than  $\{A, \dots, N\}$  would through their steps; and
- (6) Either: each of  $A, \dots, N$  would realise positive aggregate expected value if he or she were to take measures in all the cases where (1)–(5) hold;  
or: when the importance in (1) and value in (4) are used to weight the responsive steps in (3), these steps rank sufficiently highly among similarly weighted steps, for which (1)–(5) also hold, such that the steps in (3), *and* all more highly-ranked steps, could be taken by each of  $A, \dots, N$  while realising positive expected value (regarding themselves and those whose important interests are thereby targeted);

Then, in the absence of defeaters:

- (7) If (1), (2), (3a), (4a), (5), and (6), then

(7a) Either: at  $t_1$ ,  $A, \dots, N$  each have a perfect, all-things-considered duty to take responsive steps with a view to  $p$ ; or, if this individual responsiveness would be pointless unless others do likewise, to make it reasonable for the others to believe that he will take responsive steps with a view to  $p$  if he believes enough others will do likewise (and each has a perfect, all-things-considered duty to take these steps if he reasonably believes that others will do likewise); and

If (1), (2), (3b), (4b), (5b), and (6), then

---

<sup>52</sup> That is, the likelihood is proportionate to the importance of the interest(s) in (1), where more important interests are proportionate to lower likelihoods.

(7b) at  $t_1$ ,  $A, \dots, N$  each have a perfect, all-things-considered duty to take responsive steps towards there being a collective-that-can-produce- $p$ , or, if this individual responsiveness would be pointless unless enough others do likewise, to make it reasonable for the others to believe that he will take responsive steps towards there being a collective-that-can-produce- $p$  if he believes enough others will do likewise (and each will have a perfect, all-things-considered duty to take these steps if he reasonably believes that enough others will do likewise); and

- (8) If (7b), then, at  $t_2$ , once a  $\{A, \dots, N\}$  collective-that-can-produce- $p$  is formed, that collective's dependency duty will amount to a perfect, all-thing-considered duty to distribute roles that are jointly sufficient for  $p$ ; and
- (9) If (8), then at  $t_3$ , once the collective has distributed jointly  $p$  sufficient roles, each member with such a role has a perfect, pro tanto duty to act within that role with a view to  $p$ ; and
- (10) If the ranking in the second disjunct of (6) ranks multiple steps equally for some agent, and if that agent could not take all of those equally-weighted measures while realising positive expected value (regarding himself and all dependents), then that agent has a duty to take some of the tied measures, up to the threshold of positive expected value (regarding himself and all dependents), and to make it reasonable for others to believe that he will take the steps he will.

I call this “the Coordination Principle” in order to introduce “coordination” (and its cognates) as a technical term that has a disjunctive meaning. “Coordinate” means “take responsive steps with a view to a non-actual state-of-affairs in which an important interest is fulfilled, *or* take responsive steps towards there being a collective that will bear a dependency duty to take measures to realise a non-actual state-of-affairs in which an important interest is fulfilled.” The second disjunct covers both individual responsive steps taken with a view to transforming an extant collective, and individual responsive steps taken with a view to the existence of a new collective.

Condition (3) does much of the work in the Coordination Principle. Instead of deriving Ben's initial duty from his duty to perform a collective act, we derive it from a counterfactual about what would happen if Ben and a number of others each coordinated: there would either be a duty-bearing agent, or  $p$  would occur. The term “ $p$ ” can pick out any state of affairs, though—as we saw in Chapter Two—it should be described as generally as possible while mentioning all that gives it its value. This is to avoid unnecessary

proliferations of coordination duties, for example a pair of cases where  $p_1$  is described as brought about by A, B, and C, and  $p_2$  is described as brought about by A, B, and D, where the substitution of C for D makes no moral difference—but because these are different states of affairs, there are coordination duties over *both*. Rather, in such a case, p should include “brought about by A, B, and either C or D”—with the four agents *each* having duties to see to it that *this* p occurs, rather than A, B, and C having duties to coordinate around the p that refers to them *and* A, B, and D having duties to coordinate around the p that refers to them.

### ***3.3.8 Failure, Defection, and Inducing Compliance***

Allow me to elaborate on (7). If Ben’s coordination will result in p or a collective-that-can-produce-p only if enough others coordinate, does Ben have a coordination duty if too many others defect, or for whatever other reason do not coordinate? Here, there seems no point in Ben’s coordinating. Perhaps, then, he is off the hook if he reasonably believes others will not coordinate. There are two responses to this thought.

The first response is just to point out that this will be rare in practice, given what it would take for this to happen. Consider what is entailed by a duty to be responsive to the environment with a view to there being a collective or an interest-fulfilling state-of-affairs. The duty plausibly entails trying to affect the environment so as to make the collective’s (or the state-of-affairs’) existence more likely. If the environment contains agents with coordination duties, this will plausibly include convincing them that they ought to coordinate, helping them to coordinate if they look like failing, taking up the slack for their failure, and so on. This is all part of the coordination duty. (Of course, if others are culpable for their failure to coordinate, then they may well have additional duties to compensate those whose coordination was more expensive than it otherwise would have been. But these duties do not immediately fall out of considerations of dependence—the considerations with which I am concerned.) Ben is off the hook only if he has tried to coordinate in all these ways.

The second reply to the worry about others’ failure evokes the second disjuncts of conditions (7a) and (7b). Imagine one last version of the beachgoers. One person is drowning and five others are sunbathing. If they coordinated in the right way (either forming a collective or working towards the state-of-affairs directly), they could easily save the drowning person. However, Jon will lazily refrain from coordinating even if the others coordinate. Additionally, individual coordinating acts are dangerous to bystanders unless Jon also coordinates (the boat will go out of control, suppose). It’s impossible for the

others to do Jon's bit for him, as well as their own bits. Finally, acts of convincing or coercing Jon to coordinate are certain not to succeed, or will be costly (they will cause Jon to go into a rage, suppose). The swimmer drowns.

If Laura, Ben, Stan, and Julie each know that Jon will not coordinate, and if they do not coordinate for this reason but would have coordinated otherwise, then we want to capture the thought that they didn't violate any duties, while Jon did. We can capture this by saying, following roughly the proposal given by Goodin (2012), that each individual has a duty to make the world such that it is reasonable for the others to believe that he will coordinate if he believes others will do likewise. This "do likewise" locution has a self-referential function, such that it should be reasonable for each individual to believe: "you'll coordinate if you believe that I'll coordinate." Once all have signalled this, then each will reasonably believe "you'll coordinate" (since each will believe that the condition for the others' coordination—namely, everyone else's conditional willingness—is met). As Goodin (2012, 24) puts it, the duty of each is to say to the others: "I will if you will" and "I will if (you will if I will)." (Fine (2012) gives a similarly iterative analysis of joint intentions.)

This is captured in the second disjuncts of conditions (7a) and (7b). That is, if coordinating on your own would be pointless, then there is a duty to give some evidence to others that you are willing if they are willing if you are willing if they are willing—and so on. The duty to actually coordinate then does not come into existence until and unless each individual reasonably believes this of the others. Thus the vast majority of coordinations—all of those that require more than one person in order not to be pointless—will begin with individuals giving evidence of their conditional willingness.

### ***3.3.9 Too Many Coordination Duties?***

Suppose the knowledgeable beachgoers are responsive with a view to p, rather than responsive with a view to forming a collective that can bring about p. Their responsiveness succeeds—they are knowledgeable, after all—but it would have been a little less costly for them to take steps to form a collective, rather than for them to be responsive with a view to p directly. In this case, they are sufficiently well-placed both to take responsive steps towards there being a collective-that-can-produce-p, and to take responsive steps with a view to p. That is, they meet both (3a) and (3b).

My account may appear to say that their duties are to take responsive steps towards there being a collective that can rescue the swimmer, *and* to take responsive steps towards the non-drowning outcome. But if this is what my account says, my account must be wrong. According to common sense, if the best way to produce p (i.e. the way with highest

expected value for those concerned) is to coordinate with others in forming a collective that can produce  $p$ , then that is what agents should do. And this is what they should do, even if they *could* have coordinated to bring about  $p$  without a collective, though with lower expected value for those concerned.

My account agrees. It never produces duties *both* to be responsive with a view to the outcome *and* to be responsive with a view to forming a collective that has a dependency duty. This is because it is impossible to meet all of (3a), (4a), (5a), (3b), (4b), and (5b). (4a) and (4b) say that responsiveness and reforming/forming a collective, respectively, would result in positive expected value regarding those concerned. Suppose that the knowledgeable beachgoers' forming a collective would be more likely to produce  $p$  than their being responsive with a view to  $p$ . In that case, their being responsive would realise too high a cost for the swimmer—the cost of being rescued by a collective—for them to have a duty to be responsive. Condition (4b) would not be met, thus they would incur no responsiveness duty.<sup>53</sup>

## 3.4 Distributing Collectives' Duties

### 3.4.1 Distributing Collectives' Duties to Members

So far in this chapter, we have seen that collectives can bear dependency duties and that individuals can have dependence-based duties either to work towards fulfilling important interests without a collective or to work towards there being a collective that bears a dependency duty. That is, the Coordination Principle has two branches: one in which the individuals work towards an interest-fulfilling state-of-affairs without forming a collective; one in which individuals work towards a state-of-affairs in which a collective has been formed or reformed. Call the first of these the “collectivising” branch; and the second the “mere responsiveness” branch. Both of the branches are important.

Nonetheless, the collectivising branch has additional steps in it, which the mere responsiveness branch does not include. Both branches include a step at which individuals must act responsively to produce a state-of-affairs (either a state-of-affairs in which a

---

<sup>53</sup> There are some questions about how to assess the relevant value of one individual's coordinating actions, that is, of whether condition (4) of the Coordination Principle holds. This is especially difficult in cases where one individual's coordination would not make any perceptible difference to the results of the total coordination effort. Derek Parfit (1984, 76) famously gives an example where 1000 people can each contribute one pint of water to a tank that will be distributed among 1000 desperately thirsty people. Because my pint will only contribute one-thousandth of a pint to each person, *my* coordination seemingly makes no perceptible difference, so might seem to have low, or even non-positive, expected value regarding myself and the thirsty people (say, because I could drink the pint myself). I unfortunately lack space to address these issues. (For discussion, see Parfit 1984, 75–82; Otsuka 1991; Cullity 2000.)

collective is formed or reformed, or a state-of-affairs in which important interests are fulfilled). But the collectivising branch also contains a step at which the newly-created state-of-affairs (in which a collective has been formed or reformed) generates a duty for the newly formed or reformed collective. Additionally, the collectivising branch contains a step at which *the collective* must discharge that duty. The “collectivising” branch thus requires additional philosophical work to flesh out exactly how the branch operates, as it contains these extra steps at which the *collective as such* bears a duty and then discharges it. The remainder of the chapter explains the operationalization of these extra steps for the collectivising branch. These steps also apply to the straightforward cases where a collective acquires a dependency duty without individuals first having coordination duties—so the present discussion will shed light on those cases, too.

In the first of the steps, the *collective* bears a duty. Collectives can discharge duties only if their members act in the right way. So it makes sense to ask what the collective’s duty implies for the collective’s members. I will work with the following:

When a collective has a duty to see to it that X,<sup>54</sup> then each member has a duty to act (as necessary) within their role to employ the decision-making procedure to distribute roles to members in a way that: if enough members acted within their roles with a view to seeing to it that X (potentially including cajoling, coercing, and covering for others), then that would be sufficient for it being the case that X in a high proportion of likely futures. Once these X-sufficient roles are distributed, each member has a duty to act within their role with a view to seeing to it that X.

To see how this general form of collectives’ duties applies to real-world collectives, suppose the UK has a duty to accommodate N number of refugees. What duties does this entail for members of the UK?<sup>55</sup> Plausibly, the head of the department of immigration has a duty to, acting within her role, construct a task force charged with setting up facilities and procedures for processing the refugees. She also has a duty to send other UK government employees to refugee sender countries, where those government employees have duties to perform their roles of setting up refugee application and processing facilities. The Home

---

<sup>54</sup> X might, for example, be that the collective takes a measure to fulfil an important interest. The “see to it that” locution follows Pettit and Goodin (1986, 654), who assume that duties take the form “A is called upon to see, so far as possible, that-*p*,” where *p* might be that A “does something; that another specified individual does something; that something is done, no matter by whom; or that something simply is the case.” That is, the formulation allows for the full range of agent-relative and action-orientated (rather than outcome-orientated) duties.

<sup>55</sup> It is contentious who counts as a member of the state for this purpose. I address this in Chapter Six.

Secretary has a duty to oversee all of this (and the Prime Minister to check that the Home Secretary is doing so). Perhaps the UK voting public has a duty to hold the government to account in some way. So we can put “N refugees are accommodated” for “X” in the general formulation of collective duties’ distribution to members, and we get:

each member has a duty to act within their role to use the decision-making procedure to distribute roles to members in a way that, if each member acted within their role with a view to seeing to it that N refugees are accommodated, that would be sufficient for N number of refugees being accommodated. Once these roles are distributed, each member has a duty to act within their role with a view to seeing to it that N refugees are accommodated.

Note that, once the roles have been distributed, a member’s duty is *not* just a duty to perform her role, that is, to perform a specific action. It is rather a duty to *act within* her role to see to it that X. This might require using her role in ways other than those intended by the role-distributors, for example if she sees that the role distributors got things wrong in some way. Any measure a member can take that is *consistent* with the explicit content of her role, and that she is able to take *in virtue of* her role, counts as her “acting within” her role. For example, a low-level public servant’s duty to act within his role with a view to seeing to it that N refugees are accommodated might require not just (say) performing his mandated role of entering refugee data into the database, but also using his physical proximity to other public servants to motivate them to perform their roles, and maybe even performing their roles for them if they fail. These are both things that he is able to do *in virtue of*, and consistently with, the role that he explicitly has (see fn. 45).

One might object to this analysis, saying that collectives can fulfil their duties without roles being distributed to members. For example, a collective might make a contract with a non-member agent, according to which the non-member agent will pursue the good in question. So, suppose the UK makes an agreement with France that France will accommodate all the refugees. After all, the duty was just to see to it that they are accommodated—not to see to it that they are accommodated *in the UK*. If both states fulfil the terms of the agreement, we might be tempted to say that the UK “discharges its duty through” France, and that France in some sense “acts for” or “on behalf of” the UK, when it accommodates the refugees.

But notice that the UK acts when it proposes the agreement with France, negotiates the terms, and fulfils its obligations under the agreement. And these things are all done by members. Moreover, members do the work of using the decision-making

procedure to distribute roles to members that are sufficient for proposing, negotiating, and fulfilling the agreement. It is these actions of *members* that serve to fulfil the UK's duty to see to it that the refugees are accommodated—even though it does not do the on-the-ground work of accommodating them. France does not fulfil the UK's duty when it accommodates the refugees. Rather, it fulfils its own, contract-based, duty.<sup>56</sup> Thus all collective duties are fulfilled by *members* using the procedure to distribute roles sufficient for achieving an aim, and then using their roles insofar as they can to achieve that aim.

### ***3.4.2 Collectives' Duties are Made up of Individual Duties***

I have used the language of collectives' duties “entailing” members' ones. This appears agnostic on whether the collective's duty is reducible to the members' duties or not. But notice that once all the individual duties are spelled out, it is at least plausible that the collective's duty is nothing more than these individual duties to use the procedure to make the decision and distribute roles, and then to act within those roles as appropriate.

I said earlier that a collective's decision-making procedure is operationally distinct from members' decision-making procedures, and that collectives' capacities are not reducible to members' capacities—though both the collective's agency and its capacities are reducible to *facts* (in addition to the agency- and capacity-facts) about members. Thus the collective's *agency*—its decision-making procedure—is distinct from the sum of its members' individual *agencies*; and its *capacities*—what is likely to be achieved if that procedure is operated—is not reducible to the sum of its members' *capacities*. The agency and capacities are reducible only to a large number of complicated facts about member' beliefs, intentions, and commitments. These complicated facts determine the structure of the collective.

But can we give a straightforwardly reductive analysis of collectives' *duties*—the moral reasons that bear upon them on the basis of (inter alia) their agency and capacities? Yes. A collective's duties derive from (inter alia) the *collective's* agency and capacities, and they are duties to do something that the individuals cannot do alone. Yet the duties reduce to a set of individual duties to use the procedure in a certain way. That is, collectives' duties are nothing more than a conjunction of individuals' duties. The duties to actually use the distinctive decision-making procedure, or to engage in multilateralism, are not held by the

---

<sup>56</sup> It is an interesting question whether this one-off contract would make the two collectives a “super-collective.” In a very minimal sense, it would. But the point remains that the *French* action of taking in the refugees would not be what discharged the *UK's* duty to see to it that they were take in somewhere. The France-UK super-collective might have thereby discharged a duty that accrued to it, but this is a separate issue.

collective. This is because the *collective itself* does not use the decision-making procedure or engage in multilateralism. Only members do these things, so the duties to do these things must be duties of members themselves. A collective duty to use the procedure to produce multilateralism would be merely a conjunction of individual duties to do so.

Yet crucially, these individuals' duties can only exist and can only be fulfilled within the collective. Thus the collective's structure (itself reducible to individual commitments, beliefs, etc.) is an important part of the explanation and content of the individual duties to which the collective duty reduces. The collective is a framework—a way of structuring individuals—that *makes it possible* for individuals to use the distinctive decision-making procedure and that makes it likely that they will behave multilaterally.

Collective duties are a set of individual duties arranged in a certain way in a certain context. Specifically, they are arranged such that each duty depends on the assumption that the other duties will be fulfilled (i.e. on the assumption of multilateralism; the assumption that members will try), in a context where there is a collective decision-making procedure (i.e. on the assumption of collective agency). Thus collective agency and capacity are important background conditions for the individual duties that make up a “collective” duty. But there is no collective duty over and above these individual duties.

This is consistent with what I said in discussing collectives' capacities: if the collective is best-placed to fulfil an important interest and its members are not, then the collective can have a dependency duty while its members (taken individually) do not. If no member alone is best-placed, then none alone can have a dependency duty. The duty *to fulfil the important interest* is held by the collective, partly in virtue of its capacity to produce multilateralism, but this duty reduces to a number of more specific individual duties to *do what they can within their role* for that fulfilment. The collective's duty to assist is a shorthand way of referring to a set of interdependent individual duties that require a collective decision-making procedure—a certain structure of individual commitments, beliefs, and so on—in order to be fulfilled.

### **3.4.3 Collectives' Diachronic Duties and Capacities**

This reductive individualism applies to collectives' diachronic duties as well. Suppose the UK cannot accommodate the refugees now, because it doesn't have the necessary housing. But it has the ability to build the housing. It therefore cannot avoid having a duty to accommodate the refugees by citing its current incapacity, assuming the accommodating will have the same value if it occurs later on (after the housing has been built). The UK then has a diachronic duty: to build the housing, and then accommodate the refugees.

What are the implications of this for members? Here, we can simply insert “houses being built” as “X” in the general formulation for collectives’ duties’ distribution to members. Then, once the houses are built, the UK gets a new duty, where “N refugees being accommodated” gets inserted as “X.” Thus collectives, like individuals, can have duties to transform themselves and perform certain actions once transformed. These duties arise in four steps. First, members having duties to use the procedure to distribute roles to members that are sufficient for the collective to be transformed. Second, members having duties to act within those roles in such a way that the collective *is* transformed. Third, members having duties to use the procedure to distribute new roles to members that are sufficient for certain collective actions. Fourth and finally, members having duties to act within those new roles in such a way that those collective actions are performed. These are all individuals’ duties, and there are no more diachronic duties than these.

#### ***3.4.4 Failure of Collectives and Responsive Individuals***

While this thesis is primarily concerned with *prospective* moral judgments, it will be useful to say something about how my account of collective duties and coordination duties deals with *failures* of collectives and responsive individuals. This will prove relevant for the discussion of R2P, where we will find many normative claims being justified partly on the basis of some (collective or non-collective) group’s failure to discharge a duty. Distinguishing types of group failure also allows us to further consider the implications of collectives’ duties for their members, by seeing how collective culpability is reducible to individual culpability. There are at least three ways in which groups might fail to fulfil their all-things-considered duties. Two entail culpability and one does not.

First and most obviously, failure might result from negligence or malice. In collectives, one or more member might negligently or maliciously fail to act within their role in the collective to do what they can to see to it that the collective  $\varphi$ —either by failing to use the procedure to distribute roles that are jointly sufficient for  $\varphi$ , or by failing to act within one such role to do what they can to see to it that  $\varphi$ . This results in culpability for all and only those members who fail in this way. The same goes for groups of responsive individuals: if some individuals negligently or maliciously fail to act responsively in the appropriate way, then all and only those individuals are culpable. For example, if some individuals signal conditional willingness while others do not, then the latter are culpable but the former are not.

The second type of failure occurs only in collectives, and not in groups of responsive individuals. This type of failure occurs when those who set up the collective, or

who distribute roles, do not set up the collective or distribute roles in ways that safeguards against individuals' negligence or malice, resulting in collective failure. For example, they might have failed to distribute "back-up" roles, or failed to install sufficient checks and balances. They might have failed to recruit enough members to fulfil all of the collective's duties, or failed to distribute roles evenly enough among members. These people have not done all that *they* could to see to it that the collective fulfilled its duties. If this is the case, then the people who failed to set up the collective decision-making procedure adequately, or failed adequately to distribute roles, may be culpable.

Third and most complicatedly, individuals' duties might be *undermined* by others' actions or inactions. This can happen both in collectives and in merely responsive groups. For example, in the Laura-led rescue team, suppose Stan reasonably believes that even if he does his duty by starting the boat, Jon and Julie will not jump in the boat and drive it to the swimmer. Jon and Julie are terrified of water, and are certain not to go anywhere near it. If Stan acts within his role by starting the boat, then (he reasonably believes) this will be futile, as no-one except Jon and Julie knows how to drive motorboats. Here, Stan's duty is *undermined* by the fact that he reasonably believes fulfilling it would be *pointless*.

Duties can be pointless in two ways. In the way I just described, Stan reasonably believes that the collective aim to which his duty is instrumental—that the swimmer is rescued—*would not* be achieved even if he does his duty. In this case, his duty is *futile*. Stan's duty would also be pointless if he reasonably believed, for example, that Ben was going to try to start the boat whether Stan tried to or not. In this case, Stan reasonably believes that the collective's aim *would* be achieved whether or not Stan does his duty. His duty is *superfluous*. Futility and superfluity undermine individuals' duties.<sup>57</sup> The story is the same for mere responsiveness. If individuals have duties to be responsive with a view to a state-of-affairs in which important interests are fulfilled, then that duty can be undermined if they reasonably believe that their own acts of responsiveness would be futile or superfluous.

Suppose, though, that Stan's reasonable beliefs are false. In the futility case, Jon and Julie actually would have jumped in the boat if he'd started it. In the superfluity case, Ben was not actually going to try to start the boat.

---

<sup>57</sup> Lawford-Smith (2012) makes a similar point, but denies that a reasonable belief in one's role's futility *undermines* one's duty to do one's role. Rather, each member's duty just is "do your role unless you reasonably believe doing so would be pointless." This gives counter-intuitive results about people doing their duty by living in a world that happens to be some way rather than another—e.g., doing their duty just by reasonably believing that others will defect.

In the collective agent version of these cases, there is collective failure—no one starts the boat in time to save the swimmer—but there is no culpability.<sup>58</sup> We might try to find Laura culpable for not distributing the roles in a way that was sensitive to Stan’s reasonable beliefs (a failure of type two), but assuming she did not know about his beliefs, it seems we cannot do that. It is unfortunate that the collective failed in the rescue, but there is no justifiable blame.

In the mere responsiveness version of these cases, it is simply unfortunate that the important interest is not fulfilled. But neither Stan, nor Jon, nor Julie, nor Ben have done anything wrong. If Stan reasonably but falsely believes that his responsiveness would be futile or superfluous, then there is unfortunateness, but there is no collective failure (after all, there is no collective agent to have failed), and no culpability.

### 3.5 Conclusion

The distinction between random aggregates, responsive individuals, and collectives is important for allowing us to make the right judgments about the capacities, actions, and duties that exist in all kinds of (broadly speaking) “group” contexts. But arising out of these distinctions, we found a problem particular to dependency duties: what do we say when the “agent” in the dependence relationship is not really an agent, but a mere random aggregate? Or when the “agent” is a set of responsive individuals, who must somehow form or reform a collective agent in order to bear and discharge the dependency duty that we know must be in the vicinity? To solve this problem, we need the Coordination Principle, which serves as an important companion to the Dependency Principle. With my accounts of collective duties and coordination duties in hand, we are able to account for a very wide range of duties that intuitively arise out of circumstances in which the fulfilment of important interests depends on the actions of individuals taken together.

This concludes Part I: Theory, which has established a sense of “dependence” that is duty-generating. This dependence consists in having an unfulfilled important interest, where either (i) some agent is best-placed to fulfil the interest, or (ii) some group of agents is best-placed to fulfil the interest, either through (a) mutual responsiveness with a view to fulfilling the interest, or through (b) mutual responsiveness with a view to forming a

---

<sup>58</sup> Here again my account disagrees with Lawford-Smith’s (2012), according to which in these kinds of cases the collective is *culpable* even though no member is. For Lawford-Smith, the collective is culpable simply because there has been collective failure, even though every member did what they individually should. I see no reason to directly infer collective culpability from collective failure, especially given my reductive analysis of collective duties. (This is not to say there might not be some social utility in having a law that under which property is taken from the collective—but such laws might not perfectly track moral facts about culpability.)

collective that would then be best-placed to fulfil the interest. “Best-placed” applies, roughly, to that agent (or group of agents) that is sufficiently likely to fulfil an important interest if they take measures to do so (or is sufficiently likely to create an agent that will fulfil an important interest, in cases of type (ii)(b)), where the measures would have positive expected value (regarding agent and dependent) if taken in this instance, where positive expected aggregate value would be realised (regarding agent and dependent) if like measures were taken in like instances, and where the agent’s (or set’s) measures have no less expected value (regarding agent and dependent) than the measures of any other agent (set) who would discharge a duty if they had one.

My two principles, which state the precise conditions for the dependence-based duties of agents and groups of agents, might be fruitfully applied to all sorts of philosophical and practical problems. In Parts II and III, I turn to focus on two specific problems, each situated within a broader realm of enquiry: in interpersonal ethics, the problem of finding a unifying, precisifying explanation for the ethics of care; and in international ethics, the problem of finding a unifying, precisifying explanation for the Responsibility to Protect doctrine.

**PART II – INTERPERSONAL ETHICS:  
THE CASE OF CARE ETHICS**

## Chapter Four: **The Claims of Care Ethics**

### **4.1 Introduction**

Care ethics has been somewhat side-lined by analytic moral and political philosophy. Perhaps it is because the issues of crucial concern to care ethicists—contextual decision-making, personal relationships, and the role of attitudes such as compassion in moral practice—do not fit well with the “impartiality” that is characteristic of much (though certainly not all) analytic moral and political philosophy. Perhaps this is because care ethicists are sometimes more concerned with applying the theory to decision problems in applied ethics or public policy than they are with stating the theory’s central tenets (e.g., Kittay 2005; Kittay 2008; Held 2008; Koggel and Orme (eds) 2010). Whatever the reason for this historical disconnect, it might cause surprise that I now turn to fuse Part I’s theoretical framework with care ethics. Over the coming two chapters, I aim to show that this fusion is natural: Part I’s principles and care ethics have the same underlying concerns.

The basic intuition on which Part I built—the intuition that the best-placed individual or group has a duty to fulfil important interests—is, I will argue, the same intuition that underlies a plausible interpretation of care ethicists’ core claims. And my precisification of this intuition, as developed in Chapter One, does a good job of making some of care ethicists’ claims more determinate. Thus over the next two chapters, I shall argue that dependence—and the dependency and coordination duties that arise out of dependence—is a solid unifying, precisifying, and explanatory ground of a compelling version of care ethics. Far from being merely “compatible with” or “a possible extension of” my framework, care ethics is exactly where we should be looking if we want to pursue the explanatory potential of my framework. The argument will proceed by, in the present chapter, presenting what I take to be the most compelling version of care ethics’ core normative claims, and in the next, demonstrating that these claims are well- unified, precisified, and explained (justified, grounded) by Part I’s two principles.

The aim of the present chapter, then, is to develop a plausible statement of care ethics’ core normative claims. I specify that I am dealing with normative claims, because care ethicists often make descriptive claims. For example, they note that “traditional” ethical theory has problematically ignored the care work traditionally done by women in most societies, that humans are deeply relational creatures, and that humans’ embodiment makes them extremely fragile for vast swaths of their lives. Obviously, care ethicists view these empirical facts as lending support to certain normative claims. And as I consider the

normative claims, I will have recourse to assessing some of these empirical claims' ability to justify the normative claims. But my concern here is with the prescriptive conclusions of care ethics, not its descriptive premises.

In order to do this in care ethics' own terms—independently of dependence—I will put the framework of Part I to one side for the duration of this chapter. This will allow me to get into focus the care ethics that is up for explanation by this framework in Chapter Five. Developing a compelling statement of care ethicists' core normative claims is no small task, as care ethics is a diverse tradition. Space restrictions command that I paint the tradition with broad brush strokes, and not all care ethicists will agree with the details of my version of the core claims. Nonetheless, these claims constitute a related family of concerns on which there is a large consensus among those who self-identify as advocating a "care" approach to ethics (and those who self-identify as characterising the care approach). Moreover, I aim to show that a little philosophical reflection on existing statements of care ethics will lead us naturally to my version of the theory. However, reflection on the tradition will only take us so far. At the end of this chapter, we will be left with a collection of core care ethical claims that is still somewhat fragmented and indeterminate. Chapter Five's task will be to unify and precisify the loosely related and somewhat claims that this chapter has produced, by explaining it using Part I's framework.

## 4.2 Care Ethics: The General Picture

It will be helpful to begin by stating, very generally, what care ethics is. Care ethicists start by taking the phenomenology of actual ethical decision-making as crucial data for ethical theorising. They point out that, in everyday life, moral deliberation takes place in a context. When deciding what we should do in a given circumstance, we typically take account of the particularities and complexities of the relationships between the unique persons in the dilemma. We do not normally apply abstract rules or perform regimented calculi. These decision-making processes often strike us as coldly lacking in moral qualities or not quite suitable for the given situation. Rather, we remain focused on the concrete world, which contains an intricate web of valuable personal connections that call on us to make a moral response.

Early care ethicists developed these themes—contextual particularity, the uselessness of general principles, webs of relationships extended over time—from their presentation in Carol Gilligan's *In A Different Voice* (Gilligan 1982). This was a psychological study describing data in which subjects discussed moral dilemmas they were facing in their

lives.<sup>59</sup> Gilligan found that many subjects did not appeal to general principles or make categorical assertions about right and wrong. They instead used “the responsibility conception,” which “focuses on the limitations of any particular resolution and describes the conflicts that remain” (Gilligan 1982, 22). This ambivalence was a result of subjects’ perception that each dilemma involved many conflicting responsibilities. Paradigmatically, Gilligan described “Amy,” a subject who saw a particular dilemma as “a narrative of relationships that extends over time” in “a world that coheres through human connection rather than through systems of rules” (1982, 28–9). Gilligan asserted that this “orientation toward relationships and interdependence implies a more contextual mode of judgment and a different moral understanding” than one that focuses on universal ethical principles (1982, 29).

Generally speaking, care ethicists claim that responsibilities derive directly from relationships between particular people, rather than from abstract rules and principles; that deliberation should be empathy-based rather than duty- or principle-based; that personal relationships have a moral value that is often overlooked by other theories; that at least some responsibilities aim at fulfilling the particular needs of vulnerable persons (including their need for empowerment), rather than the universal rights of rational agents; and that morality demands not just one-off acts, but also certain ongoing patterns of interactions with others and certain general attitudes and dispositions. Most importantly, care ethicists claim that morality demands actions and attitudes of *care*, in addition to those of respect, non-interference, and tit-for-tat reciprocity (which care ethicists generally see as over-emphasised in other ethical theories).

Not all care ethicists hold all of these views, different theorists define them differently, and different theorists emphasise different ones. Nonetheless, this loose cluster of claims gives us some sense of what care ethicists believe matters morally. It is difficult to be more precise, as there is no generally agreed-upon statement of what care ethics *is*—providing a more precise statement is the aim of this chapter.

It is equally hard to specify exactly what care ethics is *not*. From early in its history, care ethics was contrasted to “traditional” (e.g., utilitarian and Kantian) approaches to moral theorising (e.g. Gilligan 1982; Noddings 1984; Ruddick 1980; Ruddick 1989). It is sometimes unclear exactly which other theories or theorists are being opposed here: the

---

<sup>59</sup> Gilligan’s research focused on female participants. Joan Tronto (1987) and Diana Romain (1992) convincingly separated the ethics of care from “women’s ethics.” I will put to one side care ethics’ association with feminism and “the feminine,” on the assumption care ethics’ status as a cause, effect, help, or hindrance to the various (internally contested) goals of feminism/s can be separated from care ethics’ commitments as an ethical theory. This assumption is not uncontentious (Held (2006, 20) rejects it), but unfortunately a full feminist analysis is beyond the scope of this thesis.

derisive term “traditional” is attached by those writing on care ethics to the views that caring is not morally required (Engster 2005, 57), that morality is impartial through and through (Driver 2005, 183, though Driver defends consequentialism against this charge), that all moral demands are demands of justice (Held 2006, 17), that infants do not contribute to the moral value of parent-child relationships (Noddings 1999, 36), and that humans are autonomous and independent (S.C. Miller 2005, 140).

Of course, many contemporary self-described utilitarians and Kantians would reject some or all of the claims attributed above to “traditional” morality. Nonetheless, these claims give us some initial sense of what care ethics is not—though most care ethicists recognise the limitations of drawing a sharp line between care ethics and “traditional” moral theory, or at least agree that both are important. In particular, most care ethicists now accept that care must be in some way fused or combined with its early opponent, justice, where “justice” can be understood as primarily concerned with liberty and reciprocity (Kittay 1996, 232; Kittay 1999; Held 2004, 65, 68; Held 2006, 15–17; Robinson 1999, 23ff; Ruddick 1998; Tronto 1987, 167).

Care ethicists task themselves with pointing out important ethical truths that have been neglected by other theories. But they do not claim that other theories get nothing right. Non-care ethical considerations might be entirely appropriate in some circumstances. It is likely that some combination of care and non-care concerns is appropriate in all circumstances. This is important: as I read them, care ethicists do not take themselves to be giving an account of the *whole* of morality (Daniel Engster (2007, 61–2) and Virginia Held (2004, 65, 68) seem to agree). They rather take themselves to be pointing to an interrelated set of concerns that have hitherto been inadequately appreciated.

The aim of the rest of this chapter is to make sense of the care ethical family of concerns, in a manner that is sensitive to the internal logic of the literature. I will focus on four central features of care ethics: scepticism about principles; special valuation of personal relationships; endorsement of caring attitudes; and endorsement of caring actions. Along the way, there will be reason to improve upon the most generic versions of these claims, and thus to stake out a particular conception of them within the care ethics literature. We will thus arrive at a set of care ethicists’ core claims, which will be up for unification, precisification, and explanation in Chapter Five.

## 4.3 Scepticism about Principles

### 4.3.1 *Deliberation and Rightness*

Care ethicists point out that, when actually deliberating about what we morally ought to do in some concrete scenario, we generally lack recourse to general principles. Rather, we consider concrete, particular others in complex webs of relationships. Because of their complexities, our relationships with particular others (and those particular others themselves) are an irreducible part of moral justification, deliberation, and practice. Principles—understood as conditionals with an imperative consequent—are at best insufficient, and at worst distortive, for proper moral justification, deliberation, and practice. This thought is expressed differently by different theorists, but the general idea is that a full and accurate specification of the moral reasons to perform an action,  $\varphi$ , in a context, C, will include so much detail about C that none of the reasons will apply to other contexts. So, we cannot generalise beyond C if we are to explain why the moral reasons to  $\varphi$  in C are (or are not) weighty enough to generate a duty (even a *pro tanto* one) to  $\varphi$ .

This view has a close cousin in the particularism made famous by Jonathan Dancy (2004; see also Hooker and Little (eds) 2000), according to which a reason can favour  $\varphi$ -ing in context C, and disfavour performing an action of  $\varphi$ 's type in context C\*. This is arguably a different particularism to that of care ethics: for care ethicists, the categorisation of actions into types, and the idea that the reason is the “same” between contexts, is already too general and abstract. For them, the reason is unique to this particular person or situation, and derives directly and irreducibly from the concrete things (or people or events) in this situation. In any case, despite the possible close parallels to Dancy's view, I will here focus on the view as it is presented by care ethicists.

Paradigmatically, Nel Noddings writes “[i]n order to accept the principle, we should have to establish that human predicaments exhibit sufficient sameness, and this we cannot do without abstracting away from concrete situations those qualities that seem to reveal the sameness. In doing this, we often lose the very qualities or factors that gave rise to the moral question in the situation” (Noddings 1995, 14; see similarly Noddings 2002, 20). Held offers the example of honouring one's parents, noting that the reason a child honours his or her parent is because their particular parent is worth honouring, for reasons that can be spelled out only by describing the details of that relationship (Held, 2006, 79–80).

Noddings' and Held's claims seem to be about *justifying* a particular response to a scenario. Others make the same point, roughly, about *deliberating* in response to a scenario. For example, Joan Tronto's “contextual” morality requires more than “universalizable, impartial rules” in order to be implemented. It “may” also require “a sense of the ends of

human life, an education into virtue, a moral sense, or many of these qualities" (1993, 27). Similarly, Walker contrasts the moral "understanding" of the care perspective with the (she claims, inadequate) moral "knowledge" of perspectives that use universal rules. While understanding entails "attention, contextual and narrative appreciation, and communication in the event of moral deliberation," any claim to *knowledge* about what to do in a moral dilemma will require using abstraction to make unwarranted generalisations (Walker 1989, 19–20). When Sarah Clark Miller characterises care ethics, she lists one of its four core features as "particularity," which she parses as "tending to [others] in their particularity, responding to them not as abstract 'moral patients' or 'subjects', but rather as unique individuals with distinctive life stories and circumstances" (S.C. Miller 2005, 139).

There is certainly something to these claims, particularly regarding deliberation. If we went through life trying to figure out whether the complex antecedents of various moral principles were true, and how the principles weighed up against each other, we would never get anything done. Not only that, but we would miss out on a lot of what's valuable in life: human connection, empathy, and spontaneity, for example. Moreover, one might think that if X is a good framework for moral *deliberation*, then X is likely to at least approximate the right framework for moral *justification*. It just does not seem to us that the truth of some abstract conditional, and the fact that its antecedent is true in this situation, is what makes it the case that we should, say, tell a friend his spouse is cheating on him. The specific features of the situation—in particular, the particular people and their particular relationships—seem to directly determine what we should do.

But we should be careful to separate the question of how to perceive and deliberate, from the question of right-making. Principles—understood as conditionals of the form "if C, then morally ought  $\varphi$ "—might serve as the explanation or justification of why we should do what we should do, even if entertaining that principle would be the wrong way to do that thing. Pointing this out is, roughly, the old consequentialist trick of distinguishing between decision-making procedures and criteria of rightness, or the (more recent) Kantian trick of distinguishing primary from secondary reasons (Sidgwick 1907, ch. 5; Baron 1984).

Allowing that they might have different answers, we can address the two questions in turn: first, should principles be the focus of moral deliberation? Second, do principles serve as criteria of rightness?

### 4.3.2 *Deliberation*

Even on the first question, a sweeping claim against general principles is obviously too quick. Common sense often requires that we liken current cases to previous ones; consider what this “type” of relationship requires from moral agents, and so on. This is a cognitively cheap heuristic, a way of ensuring consistency across situations, and a way of getting an answer to hard cases where an answer is required. (“Should I tell the truth?” “Usually, yes. In this case, the details are difficult to weigh up, so go with the presumption.”) However, if care ethicists can point to an alternative and more fruitful mode of deliberation, then principles should perhaps be brushed to one side—or combined in some way with care ethicists’ alternative.

And many care ethicists do offer an alternative: roughly, sympathy. In this context, we can understand sympathy as appreciating someone else’s situation from their perspective, and being moved to help them because of what one sees from that perspective. This conception of sympathy seems to capture what care ethicists are getting at when they say, for example, that “[a]n ethic reflecting concern for dependents and those who care for them demands, first, a sense of attachment to other; second, an empathetic attention to their needs; and, third, a responsiveness to the need of another” (Kittay 1996, 236), or that by “imaginatively apprehending another’s pain as painful,” the “compassionate” person is “pained by the other’s pain, and ... acts to relieve the other’s suffering” (Ruddick 1992, 152).<sup>60</sup>

Importantly for care ethicists, in sympathising we should not view the content of others’ perspectives in general or abstract terms, as we might when applying a principle. Rather, “each is entitled to expect and to assume from the other forms of behaviour through which the other feels recognized and confirmed as a concrete, individual being with specific needs, talents and capacities” (Benhabib, 1987, 164). (While not herself a care ethicist, Seyla Benhabib’s (1987) distinction between the “generalized” and the “concrete” other is used extensively by care ethicists.)

This mode of deliberation is often contrasted with something like a principles-based method. For example, Diana T. Meyers describes the “rights” perspective, in which “deliberators regard moral problems as analogous to mathematical equations with variables to compute” (1987, 141). Noddings claims that “[i]t is not just that highly mathematicized

---

<sup>60</sup> Similarly, Noddings (1984, 24) states that “[w]hen we care, we consider the other’s point of view.” Noddings later described this consideration as “receptive attention” to “what-is-there” in the care recipient (Noddings 2002, 17). This requires engrossment in, and listening to, the one cared for (Noddings, 2002, 136-37). The conception of sympathy I briefly outlined is developed outside care ethics by Cullity (2004b). For different definitions of sympathy (and its distinction from empathy), see Prinz forthcoming; Stueber 2008; Darwall 2002.

schemes are inevitably artificial (which does not mean that they are not useful), but they tend to fix our attention on their own gamelike quality. We become absorbed in the intricacies of the game instead of the plight of real people” (Noddings 2002, 60).

While something is clearly lost in the deliberation Meyers and Noddings describe, we should not take this to mean that good deliberation has no recourse to principles, for at least three reasons. First, principles are arguably just as compatible with sympathy as they are with the overly intellectualised, impartial calculation Meyers and Noddings rightly deride. Indeed, moral and political theorists of all stripes affirm the central role of sympathy in deliberation, where that deliberation also involves principles. Most obviously, sympathy is a central decision-making tool for virtuous agents, who, in some versions of virtue ethics, also abide by general principles or “V-rules” (Hursthouse 1999, Part II). Sophisticated consequentialists claim that deliberators should go back-and-forth, to some extent and as circumstances allow, between an “indirect” (that is, sympathy-, empathy-, or disposition-based) and principle-based moral deliberation (Railton 1984; on care ethics specifically, Driver 2005). Marcia Baron (1991) and S.C. Miller (2005) argue that a Kantian basis for ethics—founded on the categorical imperative in its various formulations—is consistent with a largely sympathetic or other-focused (rather than rule-focused) approach to moral practice.

While I lack space to fully interrogate these author’s arguments, the combination of sympathy and principles certainly rings true to the phenomenology of much moral decision-making, in which we establish, in a given decision scenario, a narrow reflective equilibrium between general principles, on the one hand, and the results of sympathetic appreciation of others’ perspectives, on the other. This phenomenology suggests that the choice between deliberation-by-principle and care ethicists’ deliberation-by-sympathy is not dichotomous.

Second, principles are sometimes rightly decisive in the reflective equilibrium. Consider parents engaging in “tough love,” policymakers who have to make tough decisions about funding cuts, or care workers deciding how to divide their time on a particular day. Here, principles serve to constrain the effects of sympathy on deliberation. Sometimes we intentionally put sympathy to one side, precisely in order to do what is best overall. This is at least in part because engaging in sympathy—considering another’s situation from her point of view—sometimes blinds us to other morally relevant features of the decision scenario. Of course, some care ethicists might deny that taking principles as decisive is ever appropriate, or at least deny that it is “caring.” But if this is right, then care ethics seems unable to capture much of what goes on in ordinary moral reasoning.

Third, principles and sympathy are not correctly conceptualised as engaged in a tug-of-war in the reflective equilibrium: even if we were wanting to reason entirely by principles, we would still need sympathy in order to apply principles. Sympathy is part and parcel of applying principles. Psychological findings suggest that sympathy allows us access to some of the key inputs into moral principles. Among other things, sympathy produces emotions, which help us to unconsciously and immediately narrow down the extremely large set of actions open to us, so that rational procedures can operate on the narrowed set of options. Additionally, sympathy has a key role to play in enabling us to assign *values* to the options in a decision scenario—by sympathetically adopting someone’s perspective, we are able to fully understand the importance an action will have for them, and will thus be able to see what the principles dictate, given that importance.<sup>61</sup>

We should, then, be wary of rejecting deliberation by principles in favour of deliberation by sympathy, for three reasons: first, seemingly good reasoning involves a balanced compromise between principles and sympathy; second, seemingly good reasoning sometimes allows principles to override the results of sympathy; and third, the application of principles is partly constituted by the engagement of sympathy. This is not to say that we should never act solely on the basis of sympathy. As we shall see, there are cases when acting from sympathy alone produces a special kind of moral value, simply because one acted from sympathy alone. But that is not the only kind of moral value there is—there is also value in, for example, consistency and fairness. Moral theory should endorse sympathy in deliberation, but not at the absolute exclusion of all else. That is the most that care ethics can credibly claim, but it is entirely right to claim that much. On the most charitable reading, this is what the care ethicists quoted above are getting at.

### 4.3.3 Rightness

Many care ethicists accept there might be a single criterion of rightness that encompasses all of their concerns. However, they usually do so as an argumentative strategy, aimed at demonstrating that any *unconditional* principle—any principle that can produce all and only the correct care ethical normative claims—will be so general as to be deliberatively useless. Take Noddings’ unconditional principle “always act so as to establish, maintain or enhance caring relations” (2002, 30); or Tronto’s unconditional “one should care” (1993, 153). Both of these are explicitly intended by their authors to demonstrate the *uselessness* of true, general (indeed, so general as to be unconditional), and universal principles for care ethics.

---

<sup>61</sup> D. Evans 2002; for an overview of the literature see DeSousa 2010 sec. 8.

Crucially, though, these unconditional principles are viewed by their authors as true criteria of rightness. As Michael Slote puts it within his virtue-ethical version of care ethics:

The people themselves, according to the ethic of caring, are not to guide themselves by the principle that it is right to act caringly, wrong not to; rather, they are to be directly concerned with people's well-being. But that principle can still represent a valid moral standard against which their conduct and motivation can be measured by those who would wish to do so. (Slote 1999, 28)

Tronto agrees: “[t]he problem is not that care cannot be expressed as a universal imperative: one should care ... But care is distorted if we separate the principles of care—that care is necessary—from the particular practices of care in a given situation” (1993, 153).<sup>62</sup> Tronto’s point seems to be that—if it is to guide practice—a general, unconditional principle needs to be parsed into a number of highly detailed, conditional principles, each of which might only be applicable to one particular context. Here we see playing out the tension between having principles that are detailed enough to be applicable in practice, and those that are general enough to cover a range of cases.<sup>63</sup> Almost certainly, any *unconditional* principle would err too far on the latter side of this divide. But this need not hurl us to the other end of the spectrum, where nothing morally relevant unites the sources of rightness in contexts.

Indeed, other prominent care ethicists give true *and* (somewhat) useful criteria of rightness, though they clearly intend for these to be conditional principles, i.e., principles that produce normative claims only in some scenarios. Moreover, these principles are not suggested by their authors to be *the* guiding principle of all of care ethics. (This does not make them not care ethical principles, but it does make them not the *only* care ethical principle.) For example, Engster (2007, 58) gives a “principle of subsidiarity,” which states that “we should shift the actual delivery of care whenever possible to the most local and

---

<sup>62</sup> Noddings agrees: “[o]ne might suggest as a basic principle: always act so as to establish, maintain or enhance caring relations. A carer, however, does not refer to this principle when she responds to a person who addresses her. ... The behaviour of carers is well described by this principle, but their motivation arises either spontaneously (in natural caring) or through deliberate reflection on an ideal of caring that has become part of their character” (Noddings 2002, 30–1).

<sup>63</sup> Outside care ethics, Scanlon (1998, 197–202) gives an illuminating discussion of this tension. In general, “[t]here is an obvious pressure toward making principles more fine-grained, to take account of more and more specific variations in needs and circumstances. But there is also counter pressure arising from the fact that finer-grained principles will create more uncertainty and require those in other positions to gather more information in order to know what a principle gives to and requires of them” (1998, 205). See also Korsgaard 2009, 73–4.

personal levels. We should care for others whenever possible by enabling them to care for themselves.” This clearly produces normative claims only when doing these things is *possible*. And it’s unlikely that Engster sees care ethics as non-prescriptive in cases where these things are impossible. So this is probably not meant to be *the* guiding principle of care ethics. Similarly, Eva Feder Kittay outlines a “principle of social responsibility for care”: “[*t*]o each according to his or her need for care, from each according to his or her capacity for care, and such support from social institutions as to make available resources and opportunities to those providing care, so that all will be adequately attended in relations that are sustaining” (1999, 113, emphasis in original). This is a fairly general principle, yet context makes it clear that this is a (very general) principle for social and political institutions—it does not purport to say all that care ethics might say about responsibilities in, for example, personal relationships.

When it comes to rightness, then, care ethicists are not as opposed to principles as many of them suggest when they discuss deliberation. Even those who strongly oppose principles in deliberation do not deny that there might be a deliberatively impractical criterion of rightness; and some give principles that seem at least somewhat deliberatively practical.

The upshot of care ethicists’ scepticism about principles, then, seems to be this: ethical theory should positively endorse deliberation involving sympathy and direct attendance to concrete particulars. But these modes of deliberation need not be endorsed at the absolutely exclusion of principles—indeed, they will often be necessary for applying principles. And this is not to say anything about the criteria of rightness that determine which actions and attitudes we should adopt, according to which sympathy and contextually-oriented deliberation should be employed.

## 4.4 The Special Moral Value of Personal Relationships

### 4.4.1 Initial Characterisation

Care ethicists reserve a central part of their theory for personal relationships. Of course, demarcating the relevant relationships is crucial to clarifying this aspect of the theory. But before doing that, it will be useful first to give a relatively broad characterisation of the relationships, which will allow us to see in more detail the kind of value they are purported to have. I will then be able to use the explicated kind of value to further clarify exactly which relationships are at issue.

To vaguely characterise the relationships, then, note that the term “personal” implies, first, personal knowledge of each relative by the others. In addition to this, these

relationships are characteristically not formally contracted, depend on a certain kind of history between participants, and are valued non-instrumentally by participants (outside care ethics, these kinds of characterisations are given by Samuel Scheffler (2001, ch. 6) and Niko Kolodny (2003, 148)). Paradigm examples include parents and children, siblings, friends, and spouses. Participants tend to take one another's interests as their own: it is good *for me* when something good happens to my relative.<sup>64</sup> Some personal relationships are entered into non-voluntarily, such as those between parents and children. (While one can choose to become a parent, one cannot choose to become the (biological, and often *de facto*) parent of *this* child.) Others are deliberately, though not contractually, formed—for example, some friendships and romantic relationships. Let us call those with whom we have personal relationships our “personal relatives.”

There are three main normative claims that care ethics see as arising from the special moral value of (at least some) personal relationships. These are (a) that these relationships should be taken as paradigms for morality generally (i.e., we should aim to take the same *kind* of attitude—sympathetic, compassionate—to non-relatives that we take to relatives, even if not the same extent);<sup>65</sup> (b) that some of the most morally important actions and attitudes aim to value, preserve, or promote these relationships;<sup>66</sup> and (c) regarding our personal relatives, at least some of the responsibilities we have regarding everyone are weightier.<sup>67</sup> Call these the claims of “relationship importance.”

---

<sup>64</sup> This seems to be part of what Kolodny (2003, 152) describes as “emotional vulnerability” to one’s relatives or relationships. This is a slightly less other-focused variant on what Kittay (1999, 51) calls the “transparent self”: “a self through whom the needs of another are discerned, a self that, when it looks to gauge its own needs, sees first the needs of another.” My variant is less other-focused in that I haven’t said that we see our relatives’ needs *prior to* seeing our own.

<sup>65</sup> Take Noddings, for example: “[t]he preferred way of relating to one another morally can be called *natural caring*. By ‘natural’ I mean a form of caring that arises more or less spontaneously out of affection or inclination.” For Noddings, moral practice should begin by considering how care is carried out in “the best homes.” and then extend this to others we encounter (2002, 29 (emphasis in original), 48). See similarly Ruddick 1989.

<sup>66</sup> Held argues that rather than taking utility as our starting assumption, “[w]e might … take it as one of our starting assumptions that creating good relations of care and concern and trust between ourselves and our children, and creating social arrangements in which children will be valued and well cared for, are more important than maximizing individual utilities.” (Held 1987, 126). Noddings argues that the “best starting point” for care ethics is to “take the caring relation as a primitive good.” This, she says, will have the implication that “all efforts to establish, maintain, or enhance such relations have moral worth.” (Noddings 1999, 3). Tronto claims that care ethics asks, rather than what is fair, “how might what has to be done in this situation best preserve and nurture the human relationships involved?” (Tronto 1993, 78).

<sup>67</sup> Kittay (1996, 234): “as potent as the bonds of association created by public agreements may be, they are not as powerful as those created by caring relationships. The latter are bonds that tie individuals together into families, kin, and other intimate relations, bonds that allow individuals at different stages of life to withstand the forces that act on them … Without practices based on an implicit principle of care [within intimate relations], human beings would either not survive or survive very poorly—and surely would not thrive.” Marilyn Friedman makes a case for strong responsibilities to personal relatives from a care ethical perspective, arguing that “[a] relationship, as it endures, is at its best if its participants each feel that they derive something special from their partner and have something special to offer their partner. … This recognition of, and responsiveness to, uniqueness itself seems to require that people be able to do special things for each other,

Common sense seems to agree with some version of relationship importance. Imagine a person who does not visit his senile mother in a rest home, despite living nearby. Suppose that visiting his mother would require him to recognise the third claim of relationship importance—that relationships give rise to particularly weighty duties—since there are many equally senile people in the home and he cannot visit them all. We think that his failure to fulfil this moral demand might indicate a general moral ineptitude. We think that he has more reason to value, preserve, or promote his relationship with his mother than his relationship with the other rest home residents. And we think his responsibilities to visit his mother are weightier than any such responsibility he might have to other residents. So, we criticise him, blame him, and think he has wronged his mother.<sup>68</sup> And we feel guilty if we act as he does. These judgments and attitudes suggest a failing—specifically, a failing of morality and not merely of prudence, politeness, or aesthetic preferences.

#### 4.4.2 The Relevant Relationships

Yet not just any personal relationship is important in the ways just outlined. Many personal relationships are abusive or disrespectful to participants, or have negative effects on third parties. This is despite the fact that they have the general characteristics mentioned above, and are intuitively cases of personal relationships. (Consider, e.g., abusive spousal relationships.) How should we further specify the relevant relationships? One option is to glance outside care ethics for inspiration. Here we might follow Kolodny, who argues that the relationship must have the right kind of history: a certain pattern of encounter, including shared experiences, where the totality of encounters has more value than the sum of the parts and where the encounters tend not to wrong anyone (2010b, 183ff.).<sup>69</sup> Or we might follow Samuel Scheffler, for whom a relationship generates special duties if the relatives have good reason to non-instrumentally value it (2001, ch. 6, esp. 103–4).<sup>70</sup>

---

things which cannot be done by others or for others" (Friedman 1991, 826, emphasis added). However, she goes onto argue that the conferral of extra or special care upon those with whom we are in relationships must be constrained by the effects this special treatment has on those outside the relationship, that partiality should not be practised if it inhibits general welfare (Friedman 1991, 830).

<sup>68</sup> Some reject these reactive attitudes. For example, Christopher Wellman (1997, 186–7) denies that one wrongs one's sister by choosing not to attend her wedding. He claims that this reveals one to have a bad character, but not to have acted wrongly. I do not share this intuition, quite apart from general uncertainty about the possibility of conceptually disentangling bad characters and wrong acts.

<sup>69</sup> Similarly, "one has reason to respond to a *history* of encounter in a way that is similar to the way that one has reason to respond to the *discrete* encounters of which it is composed, but that reflects the distinctive importance of a history shared with another person" (Kolodny 2010a, 183, emphasis in original).

<sup>70</sup> Scheffler claims that "to value one's relationship with another person non-instrumentally is, in part, to see that person's needs, interests, and desires as providing one, in contexts that may vary depending on the nature

Intuitively, these explanations don't go far enough. They seem to "bottom out" in a shallow place. They do not tell us why some histories (e.g. between friends) are more important (paradigmatic, valuable, normatively weighty) than others (e.g. between dentist and patient); or what a good reason is to non-instrumentally value a relationship. (Wellman (1997, 552–3) makes a similar point.) Kolodny (2003; 2010a, 186–191) successfully rules out racist or abusive relationships, but many intuitively non-valuable relationships remain potentially valuable on his view. Kolodny discusses "trivial interpersonal relationships that no one imagines provide reasons for partiality," such as a history of always boarding a train when another is leaving, and says simply that this relationship is trivial—not the reason why (2010a, 185). Other than providing a long disjunction of the right reasons for intrinsically valuing or the relevant kinds of histories, it's not clear how the Scheffler or Kolodny strategy would precede. And regarding Kolodny's proposal in particular, it's also not clear that *history* is what matters. Take a mother and her unborn child. An historical event (conception) is surely only a small part of the story about this relationship's value. It's the projection of the relationship into the future that seems most important.

Whichever way we describe the relevant relationships, there remains a problem: "relationships are ultimate sources of moral importance" is generally suspect. Relationships—similarly to food, shelter, and security—are not clearly morally valuable in themselves. What matters is the effect or meaning our actions and attitudes have for the *person*. (Pettit (1997, 155) makes the same point.) Perhaps the moral importance of our relatives has the relationship as its material or antecedent cause—we might reasonably refer to the relationship when justifying to outsiders our partiality to our relatives, for example, and it seems likely that the relationship is part of what motivates us to act as if relationship importance applies to it. But the relationship is not, intuitively, the final cause: that *for the sake of which* we should take the claims of moral importance to be true. Rather, we should act for the sake of the person themselves. So perhaps we should focus on the relevant kind of *relative*—a relative with the right kind of properties—rather than the relevant kind of *relationship*, when we are trying to figure out which relationships should have relationship importance apply to them by the most cogent version of care ethics.

An advocate of relationships-as-ultimately-valuable could respond that of course one should act for one's relative's sake, but that the relationship produces the *reason why* one should act for their sake. This view separates the moral justification (the relationship) from moral motivation (the relative). While the separation of justification and motivation is

---

of the relationship, with reasons for action, reasons that one would not have had in the absence of the relationship" (2010, 140–5).

a legitimate move—akin to the separation of justification and deliberation that I discussed above—I suggest that it takes us in the wrong direction if it leads us to relationships as sources of their own importance. When we consider what it is that might make relationships valuable or important, we are pushed on to facts about those relationships' effects on their participants—and, more generally, what we might call relationships' “meanings” for their participants (allowing that “meanings” might be broader than “effects,” as some sources of relationships' value might be valued non-instrumentally by participants).

To see this, consider that when we take away personal relationships' positive meanings for individuals, there seems intuitively little left as a basis for valuing, preserving, or promoting such relationships. Imagine a personal relationship that has a history characterised by love, affection, delight, and so on. Suppose that, if the relationship continued, it would have negative or neutral effects on the wellbeing or autonomy of both participants, with no positive redeeming effects on those outside the relationship, and that neither the participants nor anyone else values their staying in the relationship. The relationship has no valuable meaning to them or to anyone else. Presumably, the relationship would no longer be worth trying to emulate in other moral encounters, should no longer be valued, preserved, or promoted, and would no longer add extra weight to the general duties relatives have to one another just as human beings. Relationships' moral importance has limits, and those limits are determined by the relationship's meanings (including effects) for individuals.

I suggest that care ethicists should be concerned to limit the demands relationships can make when they negatively affect participants. Indeed, one thing that was troubling about Gilligan's research was the potentially *self-sacrificial* nature of her subjects' attitudes to their personal relationships—their willingness to let the relationship (or relative) make limitless demands on them (Calhoun 1988, 258–9). I suggest that to appropriately limit relationships' importance—and to articulate how they fail through subjugation, abuses of power, and inhibition of autonomy—we must turn to their meaning for the individuals in them.

In examining the status of personal relationships within care ethics, it is also important to note that contemporary care ethicists categorically deny that their theory applies only, or even mainly, to personal relationships. They instead emphasise the stringency of our care ethics-based obligations to persons all over the world (e.g., Engster 2007; Held 2006; Held 2008; S.C. Miller 2010; Robinson 2011; Tronto 1995; Tronto 2010). This has happened largely in response to early criticisms of care ethics as “parochial”: when

these criticisms were asserted, care ethicists roundly saw the need to extend care ethics to account more for non-personal relationships.

This shift has resulted in something of a tension within care ethics: on the one hand, personal relationships are still seen as important in the three ways outlined. On the other hand, non-personal relationships are recognised as important sources of imperatives to care. This might give the appearance that care ethics is now a deeply conflicted theory. But a compelling version of care ethics need not pick one side or the other. Simply, responsibilities within non-personal contexts have had their absolute moral importance increased by recognition of our moral obligations to those we do not personally know. This is not to say that personal relationships have had their absolute moral importance decreased. Rather, the effect has been that the *relative* valuation of personal and non-personal relationships has shifted. Contemporary care ethicists, by and large, do not believe that personal relationships deserve quite as much special attention, *relative to* non-personal relationships, as the early proponents of the theory did.<sup>71</sup> This seems to me the right way for care ethics to have gone.<sup>72</sup> And note that this re-calibrating of personal relationships' importance does not require full-blown cosmopolitanism: it just requires *some* weighty obligations to people we do not know. We should not see care ethicists as denying such obligations.

A sceptic of the coherence of this shift within care ethics might counter that moral theories are all about the *relative* weighting of different sources of moral value. You simply cannot “up-grade” non-personal relationships without at the same time “down-grading”

---

<sup>71</sup> Compare, for example, the globalising theories of Held (according to which, care ethics recommends international practices of “cultivating relations of trust, listening to the concerns of others, fostering international cooperation, and valuing interdependence” (Held 2006, 161)); or Fiona Robinson (“A critical feminist ethics of care grows out of a recognition of the role of power in constructing relations of dependence, upholding the myths of autonomy and concealing the needs and responsibilities of care. Thus, it recognizes the complex interdependence and relationality that characterize relations among states, institutions and individuals even in distant geopolitical regions” (Robinson 2011, 137)) with the earlier and much more parochial theories of Slote (“And like our belief in the virtuousness of caring more for those people we stand in certain special relations to, our belief in the virtuousness of greater concern for the good of (the people of) one’s own country seems to need no grounding in other ethical considerations; it is an attitude that makes sense to us, that seems preferable to treating all countries alike, even if, for example, we lack a consequentialistic or universalizability argument to use in further defense of it” (Slote 1998, 182)); or Noddings (“Our obligation is limited and delimited by relation ... I am not obliged to care for starving children in Africa, because there is no way for this caring to be completed [that is, received] in the other unless I abandon the [local] caring to which I am obligated. I may still choose to do something in the direction of caring, but I am not obliged to do so.” (Noddings 1995, 15)) Noddings’ theory was applied beyond personal relationships in Noddings 2002. There were also very early “globalisers” within the tradition, such as Sara Ruddick (1980), who discussed at length the Argentinean mothers of the disappeared: a group of mothers whose children “disappeared” during the Dirty War of the Argentinean military dictatorship, between 1976 and 1983. The mothers’ movement grew to be concerned with children’s suffering worldwide. Ruddick approvingly describes the broadening of the mother’s net of concern (Ruddick 1980, 123).

<sup>72</sup> There are innumerable arguments in support of the extension of moral concern to non-personal relatives. See Beitz 1979; Caney 2005; Brock and Brighouse (eds) 2005; Nussbaum 2006. For somewhat opposing views, see D. Miller 2007; Scheffler 2001; Walzer 1983.

personal relationships. Those of us who believe in care ethics' coherence can reply: yes, there has been a shift in the relative weightings of these relationships within the theory's literature; if that's the only weighting there is, then there has to be a shift in their weightings simpliciter. But the point is that this is a perfectly internally consistent move to have occurred within care ethics. I suggest that its consistency and motivation can be understood if we clarify what care ethicists now should take to determine the (absolute and relative) importance of non-personal and personal relationships.

Specifically, the importance of *any* relationship is determined by that relationship's meaning for (including effects on) the individuals in that relationship. Because our relationships to distant others (e.g., our concern to help them advance their political or socio-economic situation) has important meanings for us and for them, these non-personal relationships can also give rise to weighty obligations. However, personal relationships often have special kind of meaning for and effect on participants, which imbues them with a level of importance that renders true care ethicists' claims about their importance. Thus it will not always be that we must abandon our loved ones for more "needy" strangers, since our staying with our loved ones might have important meanings for them that only we can provide. (I discuss this further in Chapter Five.)

So we get Claim 2 of care ethics: to the extent that they have valuable meanings for (including effects on) individuals in the relationship, personal relationships ought to be (a) treated as moral paradigms, (b) valued, preserved, or promoted (as appropriate to the circumstance at hand) and (c) acknowledged as giving rise to weighty duties. If care ethics is to steer a middle course between the demands of personal relationships and the demands of non-personal relationships—as, indeed, most contemporary care ethicists want to do—then the value of personal relationships should be made conditional in this way.

## 4.5 Caring Attitudes

### 4.5.1 Caring About

Arguably the most important component of care ethics is that it calls upon agents to care. My interpretation of care ethics is going to have to say something informative about what this means. The two claims I have already discussed—regarding scepticism about principles and the value of personal relationships—suggest to us some facts about care, for example that it doesn't proceed by entertaining principles and that it occurs paradigmatically in good personal relationships. But we can say more than this, and we can say it more generally.

For many care ethicists, a central part of care is having certain states of mind. Tronto understands care as a “practice” that “involves both thought and action, that thought and action are interrelated, and that they are directed toward some end” (1993, 108). Kittay claims that care demands attachment, empathetic attention, and responsiveness (1996, 236–7). For Noddings, a caring relation requires that one’s attention and motivations are displaced by another’s circumstances, that one acts accordingly, and that the other acknowledges this (2002, 19). In Ruddick’s early writings (1980; 1989), the paradigm of care is found in “maternal” (and more generally, “parental”) thinking, in which “[i]ntellectual activities are distinguishable, but not separable from disciplines of feeling. There is a unity of reflection, judgment, and emotion. It is this unity I call ‘maternal thinking’” (1980, 348). Diemut Bubeck (1995, 129) describes care as an emotional state, activity, or both. Thus care ethics clearly calls upon agents to have certain attitudes, at least sometimes. (It also obviously calls upon them to perform certain actions. These are discussed in the next section. Held (2004, esp. 60) deliberately runs together the action and attitude of care, seeing them as mutually implicative—in contrast to Held, I will demonstrate that they each have value on their own, though sometimes both are morally demanded.)

The quotations also illustrate that exactly how to characterise these attitudes is hard to pin down. There are at least three options. One is to deny that there is any unitary concept here.<sup>73</sup> However, this ignores the intuitive distinctions we can make between attitudes and actions that are “caring” and those that are not. Our conceptual analysis *may* reveal multiple distinct concepts, but we should not for this reason refuse to attempt such analysis. A second option is to develop a concept that is a term of art. This has the advantages of clarity and precision, but it runs the risk of losing much of what’s intuitively valuable and important about care, and of ending up with a concept that is far removed from ordinary moral thought and practice. A third option is to explore the everyday use of “care,” and attempt an analysis that remains as close as possible to that ordinary language concept. Of course, the process of reflective equilibrium between various folk locutions and judgments might require that we reject some of those locutions and judgments as misguided or incorrect. The aim is to tidy up the folk concept while remaining true to its most central applications. I will take the third option, in part because care ethics started

---

<sup>73</sup> Ruddick (1998, 5) advocates something like this route, stating that “Attending to one’s child, without denial or projection, when she is bullying her playmates seems quite different from attending to a friend whose husband is dying.” Held (2004, 66) points out that “[t]he practices of care are … multiple, and some seem very different from others.” But she does assert that “all care involves attentiveness, sensitivity, and responding to needs.” Although this runs together the attitude and the action, which I will keep separate, these components are not far off my analysis.

as—and still is—a theory that aims to accord closely with everyday moral deliberation and practice.

The project, then, is to pick out, in ordinary language, the target for an analysis of “caring attitudes,” and then analyse that target. This will allow us to refine a broad claim—“caring attitudes are sometimes called for by care ethics”—to say, first, that caring attitudes are called for by care ethics, and second, what “caring attitudes” are. (It will not yet allow us to define the “sometimes.” That will have to wait until Chapter Five.)

#### 4.5.2 *Caring About: The Target*

According to the *Oxford English Dictionary*, “to care for” is “to take thought for, provide for, look after, take care of.” The first is an attitude, the latter three are actions. I will first focus on the attitude, and will come back to the actions. I assume that taking thought for something is caring *about* it. In ordinary language, to “care about” is to have an *attitude* of holding something to be important to oneself; to have one’s affects, desires, decisions, attention, and so on influenced by how one believes things are going with that object (similarly Frankfurt 1982; 1999). It entails a relation between a carer and an object—though the carer and the object might be identical (one can care about oneself). The possible objects of caring about are numerous: we can care about someone, something, some place, or some time. We can care about types or tokens. We might care about “interests,” or “individuals with interests,” in general. We might care about a type of event (“volcanic eruptions”), type of state of affairs (“poverty”), or a property (“having AIDS”). Or we might just care about a particular token—a particular individual, volcanic eruption, person with AIDS, etc.<sup>74</sup>

I have thus made my target for “caring about” much wider than that covered in Tronto’s definition, under which caring about “involves the recognition in the first place that care is necessary. It involves noting the existence of a need and making an assessment that this need should be met” (1993, 106). This is perhaps a good moral ideal for carers, but it is simply far too demanding as a definition: sometimes we care about someone

---

<sup>74</sup> Slote and Held each distinguish between two types of “caring about,” which we can call “specific” and “general” caring about. As Slote parses the distinction, specific caring about is an “intense personal caring towards people one *knows*,” while general caring about is “a general humanitarian caring or concern about people one only knows *about* (as part of a group)...” (Slote 1999, 2. See also Held 1993). Slote argues these two types of “caring about” must be integrated in any “morally decent person,” but that this integration needs not be conscious or deliberate. Rather, the morally decent person will simply “go about their lives, sometimes dealing with issues of justice, sometimes being involved in caring relationships - alternatively, sometimes acting from humanitarian concerns and sometimes acting out of concern of the perceived needs of people they know.” (1999, 3). This is similar to the type/token distinction I draw here, except I do not base the distinction on whether one *knows* the people. (Also, the type might not be a type of person, but rather a type of event or property or so on.)

without knowing exactly what's wrong, or without knowing what they need. Additionally, we can care about someone even if we think they have all they need: I do not cease caring about myself when I am content.

Caring can be positively valenced (a pro-attitude, e.g., "I care about scientific discovery, so I want it to continue") or negatively valenced (a con-attitude, e.g., "I care about human rights abuses, so I want them to discontinue"). Minimally speaking, we care about many things. For one not to care about something, one must be entirely indifferent to it ("I don't care about what we have for dinner"). In this minimal sense, caring about is binary—one either cares about something, or one does not—as well as ubiquitous and very easy.<sup>75</sup> Above the minimal threshold, caring about develops a scalar sense ("I care more about human rights abuses than I do about scientific discovery"). All of these are observations about the common sense term. But presumably care ethics doesn't call upon agents to have any and all of these attitudes—rather, it calls upon them to have the ones that are morally valuable.

#### **4.5.3 The Moral Value of Caring About**

Not all caring about is morally valuable. A sadist might have positively-valenced caring attitude to human rights abuses, for example—caring that they continue. When assessing the moral value of a given instance of caring about, we can ask several questions: is it morally valuable just that *there is* caring about? Is it morally valuable that *this person* care about something, regardless of what that thing is? Is it morally valuable that *this object* is cared about, regardless of who does the caring? Is it morally valuable that *this person* cares about *this object*? And for all of these: does the extent of caring, and whether it is pro- or con-caring, make a difference to the moral value?

I suggest that caring attitudes are like personal relationships: valuable only in proportion to their value *to persons*, including their meaning for persons, effects on persons, and so on. That value might lie in the attitude's being instrumental to a person's wellbeing, being partly constitutive of their wellbeing, or simply being a valuable attitude to them or for them, independently of their wellbeing. Thus caring has only extrinsic value—it is valuable in virtue of its relation to persons—but this doesn't mean that it only has instrumental value—that it is valued only as a means to some further end. Rather, caring

---

<sup>75</sup> Perhaps because of this, Noddings (1984) originally took caring about to be too thin, abstract, or lacking in real commitment to properly be part of the concept of care. In later writing (Noddings 1999), she conceded that caring about is part of the moral ideal of caring. However, I will argue below that we should separate our *definition* of care (which surely includes even minimal caring about) from the *moral value* of care. Caring about, in the limiting case, may not have much moral value. This perhaps captures Noddings' earlier concern.

attitudes might be extrinsically (but not instrumentally) valued as *expressions* of love, kindness, forgiveness, or so on—where these goods are intrinsically valuable to persons. The distinctions between extrinsic and intrinsic value, on the one hand, and instrumental and final value, on the other, are cross cutting. (Another example of something with non-instrumental but extrinsic value would be an ugly drawing done by your child: it is valued because of its *relation* to something else (your child), though it is valued not *as a means* to anything.) By and large, it is the subject, object, valence, and extent of caring about that determine its moral value. These determine its meaning for persons, and so its moral value.

An attitude's subject, object, valence, and extent might determine its moral value in all sorts of ways, given the multifarious sources of moral value that arguably exist. Care ethicists, though, are generally concerned with caring attitudes that have the right kind of relation to persons' *needs*. The term “needs” is used constantly in care ethicists' discussions of the core demands of the theory (e.g., Engster 2007, 48; Held 2006, 10, 39; Kittay 1999, 133, 233; S.C. Miller 2010, 141, 150; Noddings 2002, 88, 135; Ruddick 1998, 11; Tronto 1995, 103, 132, 133). There is some dissent: Alison Jaggar comments that beyond needs, “participants in caring relations also strive to delight and empower each other” (1995, 180). Tellingly, Jaggar presents this as an important *revision* of care ethics—suggesting that the tradition is to focus on needs. The idea of needs is usually not elaborated upon, but we can probably interpret it to include (at least, and perhaps more than) the most basic or vital constituents of, or means to, a decent life.

For care ethicists, then, perhaps the most cogent and literature-sensitive view of caring about's moral value is this. Morally valuable caring about has as its object or subject something that has, or that might affect something that has, the prospect of having a decent life, where the caring about is a pro-attitude to the having of a decent life. This is consistent with the proposition that we can care about things without that prospect, where that caring about has no moral value. It is also consistent with the proposition that we can care about things without that prospect, where that caring has moral value, if it is a pro-attitude to someone else's or something else's having of a decent life. And this is different from saying that morally valuable caring about is necessarily a positive response *to the object* of the caring. For example, to have a negative attitude to human rights abuses—that is, to be invested in such abuses discontinuing—is to respond positively to life prospects of beings that are affected by that object. Caring about human rights abuses in a negatively valenced way is morally valuable.

There is a question of what is included in our conception of a “decent life”—just how decent is “decent”? Care ethicists, and myself, need not take a stand in this. Instead

they (and I) can simply say that the worse someone's life is, the more they deserve to be cared about, all else being equal. Again, I will return to this issue in Chapter Five.

Yet I have still not clarified just why or how mere attitudes can make lives decent. To see this, consider their emotional component. Emotions often play a key role in the attitude of care, and can be used as a case study in the moral value of caring attitudes—though similar things could probably be said about the desires, decisions, attention, and so on that arise from or constitute caring about. Let us focus on the role of emotions in caring about a *person*, because it is mainly in these contexts that emotions figure in care ethics. There are many competing accounts of what emotions are, and I will not rehearse them all here (see e.g. Nussbaum 2003; Solomon (ed.) 2004; Griffiths 1997). I need not commit to any one account, though I will (fairly uncontroversially) assume that emotions have both cognitive and affective components.

Consider, then, an aged mother, Lesley, who needs to have her house maintained. In one scenario, Lesley's child, Barbara, does this out of a personal, deep, long-lasting attitude of care for her mother. In another scenario, a volunteer from the Salvation Army, Graham, does this out of a general imperfect duty of charity or beneficence. He cares about Lesley to some extent—he wouldn't like to see her hurt. But his care doesn't affect his emotions, decisions, desires, attention, and so on nearly as much as, or in the way that, Barbara's affects hers. In general, Barbara's care expresses a love and compassion that Graham does not have. Plausibly, Barbara's assistance has a significance to Lesley—and an objective moral value—that Graham's assistance could not possibly have. (A similar example is explored by Lawrence Blum (1980, 118 ff.), with similar conclusions about the instrumental and non-instrumental value of emotions. See similarly Oakley 1992, ch. 2; Stocker 1996, ch. 6.)

How can we explain this? First, having a caring attitude for a care recipient can be instrumentally valuable, by enabling attention to detail that generates *knowledge* of what this particular recipient needs and a *motivation* to meet those needs as well as one can. Second, if caring about has certain emotional components, then these might have final (non-instrumental) value, due to their relation to the proposition—in this case, plausibly “that Barbara loves Lesley”—that they express. Barbara's attitude could have this value despite the fact that Graham is equally disposed to recognise Lesley's needs when they arise, equally cognizant of the specificity of her needs, just as motivated to fulfil them (though motivated in a different way, i.e., duty), and equally aware of the desirability of fostering caring emotions within himself. This is not to say that Graham's actions are not caring actions. (More on caring actions in the next section.) But we would say that Barbara's

attitude is somehow preferable to Graham's—more valuable, perhaps more successful. And we would say this even if their outward behaviour were the same.

In the case just sketched, emotions are valued as an expression of love. But emotions can be valued as an expression of all kinds of things: respect, reverence, joy, awe, and so on. These are invariably valuable because they have the right kind of relationship to the person that feels them and to the object about which they are felt. We cannot assess the value of joy, for example, independently of extrinsic facts about it—such as who feels it and why. Nonetheless, these emotions are plausibly morally valuable precisely because they connect their subject with their object in the right kind of way.

One might claim that all these considerations are instrumental, in that they are all concerned with the effects that the emotion will have on some other thing—in the example, Lesley. This suggests that Barbara's emotions are valuable only if Lesley *knows* that Barbara is acting out of an emotion of care. So long as Barbara *believably acts as if* she is acting from an emotion, and produces all the outcomes that would be entailed by that emotion (including making Lesley believe that she has the right emotion), then her attitude has the same value as if she really did have it. Ordinary intuition moves us away from this conclusion. People can be wronged without knowing that they are wronged, and similarly, it seems, a carer's holding a certain attitude can be morally valuable despite it not being noticed by the care recipient. To demonstrate this, consider that it is wrong to feel joy at others' past suffering, even if one expresses this joy to no one, even if it is unlikely to make one cause suffering in the future, and so on. But pushing this line is not central to my project, or to my version of care ethics' core claims. Perhaps the most plausible version of care ethics will say that emotions are only instrumentally valuable. Still, their value—and the value of other manifestations of caring about—must be vindicated by any principle that claims to be at care ethics' conceptual core. And still, their value seems determined by their meaning for (including effects on) persons.

#### ***4.5.4 Duties to Care About***

A final problem remains for caring attitudes: can they coherently be “called for,” by care ethics or anything else? One might think not. After all, at least some attitudes—and particularly the desires, emotions, and so on that might constitute, cause, or result from those attitudes—seem not to be under our voluntary control. Assuming that “ought” implies “can,” we can only have duties over actions and attitudes that are under our voluntary control. While it's all very well to say that a particular attitude has moral value, it cannot coherently be demanded. We cannot demand from agents that which they cannot

intentionally deliver, and they cannot intentionally deliver that which is not under their voluntary control.<sup>76</sup>

Perhaps when they “call for” caring about, or discuss its moral value, care ethicists mean (or should mean) “it would be desirable, optimal, decent, or estimable if the agent cared about.” But they seem to mean more than this. They seem rather to implore, recommend, or otherwise make demands of agents. There are two ways of vindicating this idea for caring attitudes.

The first contends that these attitudes are under our voluntary control. We can exercise such control synchronically, for example by attending to the reasons we have to care about, downplaying the reasons against caring about, or simply acting as if we care about (with the aim that such actions will produce care) (Liao 2006, 424–5). But one might object that such “control” is not really control, since there is a low probability of our actually caring, even if we attend to reasons in the right way, act as if we care, and so on. Yet this re-casting of the “voluntary control” objection constitutes a substantial weakening of it. The claim that (certain) attitudes are beyond our voluntary control implied that there was *no probability* of our bringing them about if we tried. If there is at least *some probability* that we will care given that we do these things, then we can be morally compelled to do those things in order to have the caring attitude.

Even if it is impossible to exercise such synchronic control (which seems unlikely), it certainly is possible to exercise control over the long-term cultivation of dispositions and capacities to care about. There can then be moral prescriptions to perform such cultivation, when doing so would be morally valuable.<sup>77</sup> So perhaps prescriptions to have caring attitudes really amount to prescriptions about the long-term cultivation of dispositions or capacities. If we took this option, the statement “you ought to care about suffering” (say) would translate to “you ought regularly to attend to others’ suffering, do your best to ignore other demands on your attention, place yourself in environments where suffering

---

<sup>76</sup> Of course, if one denies that “ought” implies “can,” then the worry being addressed in this section will not get a foothold. Peter Vranas (2007) compellingly defends “ought” implies “can” against a wide range of objections. And the even more demanding principle—that moral “oughts” require *voluntary control* (seemingly more demanding than mere “can”)—is widespread. Kant claimed that “Love is a matter of feeling, not of willing, and I cannot love because I will to, still less because I ought to (I cannot be constrained to love); so a duty to love is an absurdity.” Henry Sidgwick (1907, 239) states that “it cannot be a strict duty to feel an emotion, so far as it is not directly within the power of the will to produce it at any given time.” Even Stocker—who explicitly denies that “ought” implies “can”—refuses to deny that blameworthiness implies voluntary control, when he says that “[i]f a person has no control over what he can or cannot do, over what he could or could not have done, in short, over his life, there might well be no wrong or blameworthy action” (1971, 316).

<sup>77</sup> As Blum puts it, “with regard to altruistic emotions, the prime moral task is not to control them but to bring them about in ourselves, to become a person who is prone to altruistic response when it is good and appropriate to do so” (1980, 192).

presents itself, remember or imagine yourself suffering...” and so on. To support this idea, consider that we might say quite ordinarily “one ought to feel guilty when one has not tried to help those one has a duty to help,” “one ought to feel compassion when seeing others’ suffering,” and so on. It seems we mean by these that “One ought to be the kind of person such that...” And the way to make oneself the kind of person such that X is to do the things just listed.

The second reply to the voluntary control objection is to deny that duties—and even praise or blame—entail voluntary control.<sup>78</sup> In ordinary language, we regularly praise people for actions they psychologically “simply had to do,” such as risking their lives to save others. We praise people for being such that they feel compassion for a deserted person, where this is independent of their cultivating a disposition to feel that compassion. If people do not feel such compassion, we hold them in contempt, do not admire them, look down upon them, dislike them, criticize them, or do not want to associate with them. Presumably, we praise them because their action or attitude *could have been different*, even though they lacked immediate voluntary control over whether it was different. Perhaps we praise or blame them for having or recognizing the reasons that they did—even though they could not voluntary control how compelling the reasons appeared to them. This confluence of considerations combines to suggest that care ethicists can coherently call for caring attitudes.

#### 4.5.5 *The Analysis*

The attitude of care comes in many different forms. Care ethics calls for (some of) the forms that have moral value. The case study of the emotional aspects of this attitude has allowed us to reflect upon, and get clear on the sources of, the value of these attitudes. We thus arrive at Claim 3: care ethics sometimes calls for morally valuable caring about. I suggest we understand Claim 3 in this way:

Care ethics sometimes calls for agents to have caring attitudes, that is, attitudes that:

- (i) have as their object something that has (or has the prospect of) a decent life, or something that might affect something that has (or has the prospect of) a decent life; and that
- (ii) are a positive response (e.g. promoting, respecting, revering) to that life (or prospect); and that
- (iii) lead the agent’s affects, desires, decisions,

---

<sup>78</sup> In the debate about responsibility for beliefs, several authors deny that voluntary control is necessary for responsibility (see Chuard and Southwood 2009).

attention, or so on to be influenced by how the agent believes things are going with the life-bearer (or prospect-bearer).

Clauses (i) and (ii) derive from the discussion of the intuitive moral value of caring attitudes. The conjunction of (i) and (ii) ensure that the attitude has moral value, according to the needs-focused theory of care ethics (where needs are those things required for a decent life). Clause (iii) ensures that the attitude is one of *caring about*, which on the common sense analysis is just non-indifference.

Some vagueness persists in this statement. It's not entirely clear exactly when (the most cogent version of) care ethics would call for these attitudes: for example, must I have *any* attitude that would constitute *any* improvement in people's lives? Surely this is too demanding. As we have seen, care ethicists have a partial answer to this in the idea of "needs," which we might rank for basicness or urgency. Yet exactly when agents are called upon under Claim 4 is a matter that most care ethicists leave open, and that remains open even once we have critically reflected on the care ethics literature. My proposal in Chapter Five will help to clear up this indeterminacy.

## 4.6 Caring Actions

### 4.6.1 Caring For: The Target

In addition to having attitudes, we can care by performing, practicing, or giving care. I will use the phrases "caring for" (as opposed to "about"), "giving care," and "taking care of" synonymously, to refer to actions of care.<sup>79</sup> It should be fairly obvious that the action and the attitude can come apart. The phrases "care giver" or "care practitioner" evoke an image of the professional care giver—the doctor, nurse, home helper, and so on. These people

---

<sup>79</sup> This goes against Tronto's (1993, 130–1) distinction between "Taking Care Of"—which "involves assuming some responsibility for the identified need and determining how to respond to it," requiring "agency and responsibility"—and "Care Giving"—which "involves the direct meeting of needs for care ... physical work, and almost always ... that care-givers come in contact with the objects of care." Tronto's distinction is vague and seems to lack normative significance. She seems to suggest that A takes care of B when A assigns herself the job of *seeing to it* that B's need is met; while A cares for B when A *does the work* of meeting B's need. That is, for A to be the caregiver, A must hold the last position in a causal chain of care. But "last place in the causal chain" is often vague. If Eman gives Jason money, which Jason then uses to buy food for himself, then Eman seems to be taking care of him rather than caring for him. Jason is giving care to himself (or is given care by his market transaction partner, depending). But what if Eman buys food for Jason and leaves it on his doorstep? Perhaps then she is merely taking care of him, while he is caring for himself, because she does not physically place the food in his mouth. In that case, what if she cooks the food and places it on the table in front of Jason? Still, is she merely "taking care" of him? I will run these categories together under the general heading of caring actions, though I acknowledge the political purpose of Tronto's distinction: to point out that giving money (paradigm "taking care of") is always insufficient to meet humans' needs.

perform caring actions (they care *for* others) even if they do not have caring attitudes (even if they do not care *about* those they care for). This accords with the OED's fourth listed definition of "care" as a verb, which makes no reference to attitudes:

- a. Charge; oversight with a view to protection, preservation, or guidance. [The dictionary gives the example: "I commit thee to the care of God."]
- b. "to look after ... ; to deal with, provide for, dispose of." [The dictionary gives the example: "Young ladies should take care of themselves.—Young ladies are delicate plants. They should take care of their health and their complexion."]

There are a few other closely related locutions, which are worth separating from the sense of "caring action directed at persons" that is at issue here. First, we might "care for" something, as in "be fond of" the thing. Thus when we say, "I don't care for the smell of methane," this means "I don't like the smell of methane." Second, there is "to take care of."<sup>80</sup> We might "take care of" something in the sense of seeing to it that the thing occurs: "I'll take care of the party." Neither of these are the sense of care I am discussing. They are instead different concepts that happen to be picked out by similar locution. (Exactly how they differ, I will get to presently.)

Like the attitude, the action "care" always entails a relation between a carer and an object. Like caring about, the carer and the object might be identical (one can care for oneself; one can care about oneself). I said that "caring about" has both a binary sense (we care about anything we are not indifferent to) and scalar sense (we care about some things much more than others). "Caring for" similarly has both a binary and a scalar sense: a doctor can care for two patients, but care for one of them *better* or *more successfully* than she cares for the other.

However, in ordinary thought and talk, "caring for" has a smaller range of possible objects than "caring about." In the sense I am interested in, one does not *care for* a type of event ("volcanic eruptions," "human rights abuses," "scientific discoveries"), or a type of state of affairs ("poverty"), or a property ("having AIDS"). We might care for (as well care about) those who are *affected* by volcanic eruptions (human rights abuses, scientific discoveries, having AIDS), but then we are not caring for these things themselves.

In addition to applying to a smaller range of objects, care as an action is not open to a "pro" and "con" reading: to care for something is always to respond positively, rather

---

<sup>80</sup> See fn. 79 on Tronto's definition of this locution. Tronto develops it as a term of art quite far away from the ordinary language use.

than negatively, to that very thing. And the minimal (binary) form of “caring for” is harder to instantiate than the minimal form of caring about. While we care about anything we are not indifferent to, caring for requires actions that are directed towards the object in a certain kind of way. This “certain kind of way” is what narrows the category of things it is possible to care for: we can care for something only if we can intend our actions towards it in that way.

Specifically, all caring actions are intentional under the description “trying to do what I believe is good for someone.” Why just “trying” and “what I believe”? Consider a child who keeps a rock wrapped up in a blanket, carries the rock around with him, asks people to be quiet when he believes the rock is sleeping, and so on. He is asked whether he is caring for the rock, and he answers affirmatively. Does he actually care for the rock? He at least intends to. He intends to look after the rock, tend to it, enable it to live well, and so on. In short, he does what he believes is in the rock’s interests.

I suggest that the boy does, in fact, care for the rock. He just does not do it very well. It is often difficult to distinguish doing something badly from not doing it at all. If I join in the department’s weekly social soccer match, but play terribly, I am still playing soccer just as long as others include me in the game and I am trying. If I get out paints and a canvas and use them to represent the cockatoo outside my window, then I am painting the cockatoo even if it is unrecognisable as such. In these cases, the actor’s intentions (along with, perhaps, social conventions) are key to determining whether one is playing soccer or painting the cockatoo. For caring, I suggest, it is *all* in the actor’s intentions. To care for someone is to do what you believe is in the interests of that thing—even if that thing, in fact, lacks interests, or even if you are incorrect about their interests.

The intention is not a terribly strict condition. The carer need not consciously entertain their intention as “doing what I believe is in the recipient’s interests” and they need not have a full-blown concept of interests. They just need a tacit belief that the action is good for the recipient in some way. Children, for example, can perfectly well care for their parents, without a hint of reflection on the fact. Moreover, fulfilling the person’s needs need not be the final intention of the carer—the care can be intended to be instrumental to some other aim. Consider a doctor who meets a patient’s needs only because he will get a pay check for doing so. He cares for the patient, despite not caring about her.

In line with these considerations, I will use the following definition of caring action:

an action is caring if and only if it is performed under the (perhaps tacit) intention of fulfilling (or going some way to fulfilling) interest(s) that the agent perceives some moral person (the recipient) to have.

Some will claim that this definition is too broad. For example, Bubeck describes care such that, by definition, it fulfils needs. She defines care as an emotional state, activity, or both, that is functional and specifically involves “the meeting of needs of one person by another where face-to-face interaction between carer and cared for is a crucial element of overall activity, and where the need is of such a nature that it cannot possibly be met by the person in need herself” (1995, 129). But there seems no reason to exclude care that occurs through non-face-to-face contact (consider telephone counselling), or to exclude the possibility of meeting someone’s needs through care even if they can meet their needs themselves (though the care might not be morally demanded). Bubeck also denies that one can care for oneself. I take her definition to be too narrow in all these ways, and similarly too narrow in its demand for efficacy. This just isn’t true to the common sense notion.

Engster’s definition is similarly narrow: “[e]verything we do directly to help individuals to meet their vital biological needs, develop or maintain their basic capabilities, and avoid or alleviate unnecessary or unwanted pain and suffering, so that they can survive, develop, and function in society” (2007, 28–9). Engster adds that this must all be done “in an attentive, responsive, and respectful manner” (2007, 31). The phrase “do directly to help” is important here: care is done with the aim of helping individuals. In focusing on intentions, Engster’s account agrees with mine. However, it seems we perfectly well care for someone if we enable them to survive, develop, and function *outside* society. And as with Bubeck, it’s not clear why Engster excludes from caring all the frivolous, non-vital, non-basic, life-enhancing things humans do for one another. Engster states that he does not want to posit a “broader” definition that “too closely associates caring with a particular liberal understanding of the good life” (2007, 27). Yet we can allow that caring occurs in ways that are particular to various conceptions of the good life, without associating it with only one such conception. Caring for a child might involve taking them to church, or not, depending on one’s conception of the good life.

My definition steers a middle course between Bubeck’s and Engster’s narrow definitions, on the one hand, and Berenice Fisher and Joan Tronto’s much-cited broad definition, on the other. For Fisher and Tronto, care is “a species activity that includes everything that we do to maintain, continue and repair our ‘world’ so that we can live in it as well as possible. That world includes our bodies, our selves, and our environment, all of which we seek to interweave in a complex, life-sustaining web” (1990, 40). That this

definition focuses on changing the world in general loses the *relational* aspect of care: it is *from* someone *to* someone. (As Ruddick puts it, “[t]he work [of caring] is constituted in and through the relation of those who give and receive care” (1998, 13–14).) Fisher and Tronto’s definition also seems to include too many actions in “care,” such that care becomes assimilated with just “doing good.” This assimilation might have been Fisher and Tronto’s intention, but I think there is a distinctive and morally valuable class of actions called “caring” that is worth distinguishing from generally “making the world a good place.” Caring is distinct in that it is directed at some being with interests, and in that it arises from the (perhaps tacit) intention of fulfilling (or going some way to fulfilling) the interest(s) the agent perceives that being to have.

#### **4.6.2 The Moral Value of Caring For**

In defining caring action, I have talked only about the intentions and beliefs of the caregiver. But when it comes to the moral value of caring action, common sense intuition tells us that their effects matter, too. Take the boy and his rock. Take “success” to mean “the fulfilment of the interests that the agent intends to fulfil.” The boy’s actions have less moral value than they would have done if they were successful. The boy’s actions are not fulfilling, or going some way to fulfilling, any interests of the rock. The rock does not *have* any interests. The boy’s actions are caring (since he believes the rock has interests, which he intends to fulfil). But we want to say that the boy’s care for the rock is less successful than it would have been if the rock had interests that he’d fulfilled.

The boy’s actions of caring for the rock might—*independently of his attitudes*—still have some moral value: the actions might serve as *good practice* for situations later when such care would be successful; there might have been some *non-zero likelihood*, at the time of action, that the action would come to be successful (thus we might value his having acted on this possibility). But success (fulfilling the recipient’s interests) plays a large part in determining an action’s moral value—probably larger than these other factors. Actual effects matter.

Of course, what constitutes an action’s moral value should be separated from the conditions under which an agent morally ought to perform the action. As I suggested in Chapter Two, duties are most plausibly affected by the *expected* value—including expected chance of success—not whether success occurred or not. Yet when we are *retrospectively assessing* the value of an action (irrespective of whether there was a duty to perform it), its actual effects loom large.

Allowing that actual effects matter for the retrospective assessment of actions, some effects will matter more than others. In Chapter Two, I pointed out that not all interests are created equal. This point is also frequently made by care ethicists, who, as we have seen, tend to discuss care's effects on persons' needs—where needs are only a sub-set of interests. While care might be directed at fulfilling any interest—however trivial—care will have more moral value if it fulfils a more vital or compelling interest (a need).

And actual effects are not the only things that matter. Intentions can have instrumental moral value in much the way attitudes can—that is, they can make us more attuned to interests, more motivated to fulfil them, and so on. Plausibly, they can also have non-instrumental moral value, as attempts to fulfil interests that they are not actually instrumental to fulfilling. (Consider the boy and his rock.) This value is extrinsic, though non-instrumental, as it depends on the intention's non-instrumental relation to a (perhaps merely perceived) interest, that is, that the intention is a certain kind of response to an interest. Note that the intentional component of “caring for” is quite separable from the mental states involved in “caring about”—we might intend to care for someone, despite not caring about them. (Consider the callous doctor—she intends to care for her patients, but does not care about them. And the intention might be valuable even if inefficacious.)

A final point concerns the role of the recipient in all this. Many care ethicists assert that care is successful *only if* the recipient responds to it positively.<sup>81</sup> While the recipient's positive response might be necessary for some interests (such as autonomy or empowerment) to be met, it should not be taken as essential to care—even to morally valuable care—across the board. For example, patients in healthcare may know that they have a need but not know that what is given to them will meet the need. Or they may have adaptive preferences that make the state of neediness appear normal and prevent them from recognizing the need at all. Or they may engage in wishful thinking, which dissuades them from believing that care is necessary to overcome the need. Or they may discount future rewards and so not appreciate the future benefits of, for example, preventative care. Suppose B is in a coma, and A makes a real attempt to meet B's physical needs every day. Even if B never awakes from the coma, and so never recognizes that A cared for B, it seems that A's actions should constitute care. (Similarly, if it's actually in B's interests to die, then A's euthanizing B because of this interest should constitute care: it all depends on A's intentions and B's interests.) Of course, it should be obvious that some back-and-forth

---

<sup>81</sup> Noddings gives three necessary conditions for a caring relation, one of which is “B *recognizes* that A cares for B” (2002, 19, emphasis added). One of Tronto's four phases of care is “Care Receiving: A has successfully cared for B when B accepts the care A has given. This requires responsiveness and communication between A and B” (1993, 106–7). See similarly Ruddick 1989, 180.

between carer and recipient, *where that is feasible*, will often increase the value of all kinds of care. And some of the most valuable caring relationships are reciprocal (think about spouses, for example). In these cases, the success of reciprocal care will often hinge upon mutual feedback about past caring actions.

The *moral value* of caring action, then, is plausibly a function of (1) how well that action fulfils the recipients' interests (where needs are weighted more heavily than other interests), and (2) the strength of the agent's intentions to fulfil the recipient's interests. (2) is separable from (1).

To sum up: an action can be an instance of *caring for*—can be a *caring action*—despite having little moral value, just as long as it has the right intentions (i.e. fulfilling someone's perceived interests). Above this threshold, a caring action can have more or less moral value, as a function of (1) and (2) above. The moral value of the action is separable from (3): the moral value of *attitudes* (caring about) that might accompany the action. However, (1), (2), and (3) can all be combined for an overall moral valuation of the relation between carer and the object of care.

#### **4.6.3 The Analysis**

The upshot of the discussion of caring for is an analysis of Claim 4: care ethics sometimes calls for morally valuable caring for. The analysis is this:

Care ethics sometimes calls for agents to perform actions that (i) are performed under the (perhaps tacit) intention of fulfilling (or going some way to fulfilling) interest/s that the agent perceives some moral person (the recipient) to have; (ii) where the strength of the demand is a complex function of the value of the intention, the likelihood that the action will fulfil the interest, and the extent to which the interest is appropriately described as a “need.”

As with Claim 3, we are left with some radical indeterminacy here. It is now clear what caring actions are, but it remains unclear when the “sometimes” is in which they are called for. Again, reflection on the internal logic of care ethics does not readily reveal a clarification of this. That will be part of the job of Chapter Five's unifying and precisifying explanation of the four claims of care ethics developed in this chapter.

## 4.7 Conclusion

This chapter intended to sort through the various normative claims that are made within care ethics, in order to develop the most coherent and convincing version of the core normative claims of care ethics. I started by discussing care ethics quite generally, before honing in on four large areas of concern, arriving at the following claims of care ethics:

Claim 1. Ethical theory should positively endorse deliberation involving sympathy and direct attendance to concrete particulars.

Claim 2. To the extent that they have valuable meanings for (including effects on) individuals in the relationship, personal relationships ought to be (a) treated as moral paradigms, (b) valued, preserved, or promoted (as appropriate to the circumstance at hand) and (c) acknowledged as giving rise to weighty duties.

Claim 3. Care ethics sometimes calls for agents to have caring attitudes, that is, attitudes that: (i) have as their object something that has (or has the prospect of) a decent life, or something that might affect something that has (or has the prospect of) a decent life; and that (ii) are a positive response (e.g. promoting, respecting, revering) to that life (or prospect); and that (iii) lead the agent's affects, desires, decisions, attention, or so on to be influenced by how the agent believes things are going with the life-bearer (or prospect-bearer).

Claim 4. Care ethics calls for agents to perform actions that (i) are performed under the (perhaps tacit) intention of fulfilling (or going some way to fulfilling) interest/s that the agent perceives some moral person (the recipient) to have; (ii) where the strength of the demand is a complex function of the value of the intention, the likelihood that the action will fulfil the interest, and the extent to which the interest is appropriately described as a “need.”

These core claims do not capture everything that is of concern to all care ethicists. And I should reiterate that they are merely the *normative* claims of care ethics. Many care ethicists hold certain empirical views that are crucial to their overall ethical outlooks. For example, many care ethicists endorse a relational view of autonomy, according to which our plans, projects, and purposes are inseparable from and hugely influenced by those around us. Many care ethicists emphasise that the world of ethical concern is constituted by complex webs of relationships between fragile, embodied human beings. With the core normative claims above on the table, it is easy to see how they arise out of a deep appreciation of these empirical claims. In this chapter I have been concerned with care

ethicists' normative claims, since these are the ones that are up for unification, precisification, and explanation in the next chapter.

These four claims capture what is distinctive about care ethics. Although moral theorists who do not call themselves care ethicists may endorse these four claims, the claims are unlikely to be the central or most important claims of their ethical theory. They are unlikely to be interested in intricately analysing actions of *care* in particular, or in vindicating sympathetic modes of deliberation in particular, as part of the central part of their theoretical edifice. It is the combination of these claims, and their status as the most important aspects of the theory, that marks out care ethics.

Most importantly for my purposes, though, notice that these four claims are only loosely related. Clearly, they're motivated by something like a concern for compassion, the personal, and the intimate. But it's not entirely clear what *one* common normative basis they might have, if any. It's not entirely clear that they're anything more than an ad hoc list of claims that are endorsed by people that have come to be called "care ethicists." Also, as we've seen, some of them remain problematically indeterminate. It is not clear exactly when care ethics calls for caring attitudes or caring attitudes, or in exactly which situations (if any) a non-sympathetic or not entirely contextually orientated kind of deliberation might be appropriate. Probably, care ethicists are imprecise about this intentionally, being sceptical of any general rule. In the next chapter, however, I argue that we can do better—that we can have a unified, precisifying explanatory basis for these various claims, that removes (at least partly) their ad-hoc and vague appearance.

## **Chapter Five:** **Unifying and Explaining Care Ethics**

### **5.1 Introduction**

We now have a statement of an initially plausible version of care ethics. But this statement—or, rather, set of four statements—remains disjointed, and some elements of it remain indeterminate. If care ethics is to be maximally appealing from an analytic point of view, there should be some unified core within this apparent disjointedness and a way of rendering the indeterminacies determinate. If these four claims can be unified, precisified, and explained by one core idea, then this will provide guidance in interpreting the claims of care ethics and applying them to issues in the real world. It would allow care ethicists to distinguish themselves from non-care ethicists and to determine whether new proposals in care ethics are true to the guiding concern. It would also allow them to appropriately constrain their theory's upshots. Care ethics would be better theory, I suggest, for having a precisifying and unifying explanation. This gives us reason to see if we can find one.

This chapter will present a positive proposal for precisifying, unifying, and explaining of the version of care ethics developed in Chapter Four. The idea is to demonstrate what a plausible, precise, unified version of care ethics might look like, by laying one such option on the table. To do this, I will bring together the Dependency and Coordination Principles from Part I and the claims of care ethics from Chapter Four. The aim is to demonstrate that dependence—and the two moral principles of dependence that I developed—provide a plausible, unified, explanatory ground for the care ethical claims.

If I am right, this will be important not only for care ethicists, but also for those of us who want to explore the full normative implications of Part I's principles. It will mean that these principles, when put into practice, plausibly demand agents to reason by empathy, acknowledge (epistemic) indeterminacy, and view the world as a complex web of relationships that deeply affect one another. That is, proponents of Part I's principles might be surprised to learn that they have good reason to endorse care ethics' claims; that by being committed to these principles, they are committed to a range of normative ethical claims that they may not have anticipated.

Of course, even if Part I's principles produce the four claims of care ethics, it will remain true that they produce other claims as well. As we saw in Part I, they produce claims about situations to which the four care ethical claims do not obviously apply (though they do not obviously not apply), such as one-off rescue cases. And in Part III, I will argue that Part I's principles have application to issues in international ethics that are

radically separate from care ethics. But the point is to show that the principles give us—amongst other things—the care ethical claims.

I start in §5.2 by giving some initial evidence, from within care ethics, for a dependence-like unified ground of the theory. §5.3–5.6 examine each of the four claims in turn. For dependence—and the dependency and coordination principles that reflect the moral importance of dependence—to be a unifying, precisifying explanation of care ethics, two things must be true regarding each of the four claims. First, the deliverances of the two principles within the four domains should align with those of the care ethical claims. (The four domains correlate to the four claims of care ethics, that is, the four domains are: sympathetic and contextually-orientated deliberation, personal relationships, caring attitudes, and caring actions.) Second, the principles should provide a good explanation of the claims' deliverances regarding these domains.

The first part of the explanation is impossible to demonstrate with certainty, because the care ethical claims do not always give a determinate answer about whether there is a care ethical duty, or under which exact circumstances their general statements about value give rise to actual value. For example, as I developed them, they say that sympathy should be endorsed (though not necessarily always), and that caring attitudes and actions are sometimes called for (but perhaps not at all times). So my task here will partly be to show that care ethicists, given their broad commitments, *should* endorse the duties the principles generate regarding sympathetic or context-orientated deliberation, personal relationships, caring attitudes, and caring actions.

The second part of explanation will appeal to intuitive judgments about reasons, to justificatory statements within the literature, and to possible rival explanations. However, I will address possible rival explanations only when they are particularly salient, because this chapter is engaged primarily in the positive project of suggesting a basis for care ethics, rather than the negative project of rejecting other bases.

## 5.2 Initial Evidence for a Dependence-like Ground

Part I developed “dependence” as a technical term. To recall, dependence consists in some person’s having an unfulfilled important interest, where either (i) some agent is best-placed to fulfil the interest, or (ii) some set of agents is best-placed to fulfil the interest, either through (a) mutual responsiveness with a view to fulfilling the interest, or through (b) mutual responsiveness with a view to forming a collective that would be best-placed to fulfil the interest. “Best-placed” applies, roughly, to the agent (or set of agents) that meets the following conditions: first, who is sufficiently likely to fulfil the important interest if

they took measures to do so (or is sufficiently likely to create an agent that will be is sufficiently likely to fulfil the important interest, in cases of type (ii)(b)); second, whose most efficacious measures would have positive expected value (regarding agent and dependent) if taken in this instance; third, who would realise positive expected aggregate value (regarding agent and dependent) if like measures were taken in like circumstances; and fourth, the agent's (or set's) measures have no less expected value (regarding agent and dependent) than the measures of any other agent (or set) who would discharge a duty if they had one. In this chapter, I will usually discuss “dependence-based duties” quite generally, referring to all cases of types (i), (ii)(a), or (ii)(b), though in a few places it will be necessary to distinguish the individual and group variants of dependence (as developed in the Dependency and Coordination Principles respectively).

Obviously, this algorithmic formulation is quite alien to care ethics. Nonetheless, there is widespread evidence within the tradition that, *if* there were a general ground for the theory, something like dependence would be it. The term “dependence” is used frequently within care ethics, to refer to a general type of relation of which the conception outlined above is one interpretation. Part of the aim of this chapter is to show that my particular, technical conception can usefully precisify the various claims of care ethics, so it should not be surprising that my particular detailed conception is not employed by care ethicists. Yet my conception is at least a candidate for grounding the theory, since the more general idea of “dependence” is rife in the care ethics literature.

There are numerous examples of this. Kittay argues that the responsibility to care is based upon “our unequal vulnerability in dependency, on our moral power to respond to other in need, and on the primacy of human relations to happiness and well-being” (1999, 113); and that politically “we need a concept of interdependence that recognizes a relation not so much of reciprocity as of nested dependencies, linking those who help and those who require help in order to give aid to those who cannot help themselves” (1996, 233). The opening paragraph of her influential book *Love's Labour* (1999) uses the idea of dependence repeatedly.<sup>82</sup>

S.C. Miller (2005, 140–41) lists “dependency” as one of four key features of care ethics, along with “interdependency,” “need,” and “particularity.” She parses “dependency”

---

<sup>82</sup> Kittay's book opens: “Dependents require care. Neither the utterly helpless newborn who must be cared for in all aspects of her life nor a frail, but functioning, elderly person who needs only assistance to carry on with her life, will survive or thrive without another who meets her basic needs. Dependency can be extensive or brief, with the extended dependency of early childhood or temporarily incapacitating illness. Dependencies may be alleviated or aggravated by cultural practices and prejudices, but given immutable facts of human development, disease, and decline, no culture that endures beyond one generation can be secure against the claims of human dependency” (Kittay 1999, 1).

as “situations of significant reliance on others that all persons undergo during the course of their lives. ... The certitude of dependency ... tells [a] story, one in which how we do or do not care for one another in our shared moments of dependence marks a matter of great moral importance.” (2005, 140). Miller claims that “[a]s finite and interdependent moral agents, we are required to respond to others’ fundamental needs” (2010, 150).

When Held characterises the “major features” of care ethics, the first feature is:

the compelling moral salience of attending to and meeting the needs of particular others for whom we take responsibility ... The ethics of care recognizes that human beings are dependent for many years of their lives, that the moral claim of those dependent on us for the care they need is pressing, and that there are highly important moral aspects in developing the relations of caring that enable human beings to live and progress. ... [t]he ethics of care stresses the moral force of the responsibility to respond to the needs of the dependent. (2006, 10)

Along similar lines, Tronto suggests that “we might assume responsibility because we recognize a need for caring, and there is no other way that the needs will be met except by our meeting it” (Tronto 1993, 132). In distinguishing care ethics from theories of liberal justice, Engster outlines care ethics thus:

Care theory ... begins with individuals already existing in society and dependent upon one another for their survival, development, and social functioning, and highlights the unchosen obligations we all have toward others by virtue of our interdependency. Because we are all born into a state of dependency and depend upon others more or less throughout our lives, all capable individuals have obligations to care for others in need regardless of our explicit or tacit consent. (2007, 7–8)

Other care ethicists emphasise need—though not necessarily needy persons’ dependence on others—as the driving normative force of care ethics: Ruddick states that “meeting the needs of other people is paradigmatic of care” (1998, 11); and Noddings (2002, 88, emphasis added) claims that “[i]t should matter to us that someone is suffering, and this mattering does not depend on some moral equality inherent in persons.”

Engster develops a consistency-based justification for care responsibilities: we must be consistent about when help is demandable for those in need. Engster develops “the

principle of consistent dependency”: “you must do unto others *as you have already had others do unto you*” (Engster 2007, 54, emphasis in original). According to Engster (2007, 48), “[i]n claiming care from others, we imply that capable human beings ought to help individuals in need when they are able to do so consistent with their other caring obligations.”<sup>83</sup> Engster’s idea is that we cannot claim care from others unless we accept the implication, and if we accept the implication then we are committed to a general principle of helping (or caring for) those in need. As it happens, all natural individuals who are capable of caring have already received care early in their lives, so they are all committed to caring for those in need. Thus, on Engster’s view, it is humans’ inevitable dependence—combined with the consistency principle—that produces obligations to care.<sup>84</sup>

The language of dependence is used by care ethicists to lay the foundations not only of individuals’ responsibilities, but also of collectives’ responsibilities. In describing her “critical feminist” version of care ethics, Robinson states that “[a] critical feminist ethics of care grows out of a recognition of the role of power in constructing relations of dependence, upholding the myths of autonomy and concealing the needs and responsibilities of care. Thus, it recognizes the complex interdependence and relationality that characterize relations among states, institutions and individuals even in distant geopolitical regions” (Robinson 2010, 137).

None of this is decisive evidence for Part I’s particular principles as underlying these theorists views, but it does suggest a certain polemic trend. Additionally, the assertion that dependence, vulnerability, and need permeate human life is an *empirical* claim that care ethicists cling tightly to. (Though we are not, of course, all dependent on others to the same extent or in the same ways. Kittay’s (1999) and Tronto’s (1993) versions of care ethics, for example, emphasise the radically unequal extent to which persons are dependent on others.) Theorists’ concern with this empirical claim would be explained if they saw these notions as grounding their normative theory.

For example, they point out that when we are young, ill, and old, we cannot live if not connected to others in the right ways. And some humans are helpless in this way all time (Kittay 1999). Moreover, most of us need others if we are to pursue our various

---

<sup>83</sup> Kant also took this line: “since our self-love cannot be separated from our need to be loved (helped in case of need) by others as well, we therefore make ourselves an end for others; and the only way this maxim can be binding is through ... our will to make others our ends as well.” (*Doctrine of Virtue*, 6: 393; quoted in Cottingham 2010, 74). S.C. Miller (2005) explicitly connects Kant’s line of reasoning to care ethics.

<sup>84</sup> Of course, one might deny that the consistency principle grounds obligations to dependents, because it makes these obligations worryingly contingent. Surely we would still have such obligations even if we had not received care when we were young. As I discussed in Chapter Two, my conception of dependence is relatively neutral among foundational theories. It is thus neutral about whether Engster’s consistency principle is the right ground of the duties he recognises we have to dependents.

conceptions of the good life. This is because, for most of us, having certain kinds of relations with others is part of the good life. And we depend on those around us to help us *form* a conception of the good life: our projects and preferences are greatly informed by those around us. These ideas lead many care ethicists to endorse a relational view of autonomy—roughly, the view that our “true” or “authentic” (i.e., autonomous) preferences depend for their content upon our relationships with others (Mackenzie and Stoljar (eds) 2000). And actually exercising our autonomy, by forming the intention to pursue those preferences, often requires having the right kinds of relationships with others, namely, ones that foster our self-worth and a conception of ourselves as an agent.

Along these lines, care ethicists characterise humans as “relational and interdependent,” not as “self-sufficient,” “self-interested,” or “independent” (as Held 2006, 13–14 puts it when characterising care ethics in general). Without the fact of human (inter)dependence, S.C. Miller claims, “the moral self becomes an unrecognizable wisp of moral abstraction” (S.C. Miller 2010, 149). For Noddings, the self just *is* a set of affective or meaningful encounters with things in the world (2002, 97ff.).<sup>85</sup> Care ethicists use our pervasive dependence on others—for both life and autonomy—to demonstrate that non-interference is not what we primarily need from others. Rather, we need care. Here again, humans’ dependence on others for various goods is emphasised by care ethicists.

Finally, although care ethicists generally eschew proclamations about generally-described duties that all agents have, there is one class of duties that many care ethicists are particularly concerned to vindicate. These are duties to “utter dependents,” which Kittay defines as meeting three conditions:

First, the dependent requires care and caring persons to meet the fundamental needs for survival and basic thriving. Second, while in the condition of dependency, the dependent is unable to reciprocate the benefits. And, third, the intervention of another is crucial to ensure the needs of the dependent are met and that the interests of the dependent are recognized in a social context. (1996, 220)

This category includes infants, children, and mentally disabled adults. Many care ethicists view vindicating duties to utter dependents as a basic desideratum of a moral or political theory. This is evidenced by the fact that writers use care ethics to critique Rawlsian

---

<sup>85</sup> In this, care ethicists follow Baier, who argues that we are all “second persons” who inherit our values and beliefs from those around us, and construct our identities out of them (1981), and Nedelsky, who argues that a shared conceptual apparatus is necessary for understanding the attitudes, character traits, desires, and needs of oneself and others (1989, 11).

liberalism for failing to adequately capture such duties (e.g., Bhandary 2010; Kittay 1996; Noddings 2002, ch. 4). Of course, care ethicists' endorsement of duties to *utter* dependents does not establish their commitment to dependence-based duties more generally, or to dependence as the basis of (say) duties within personal relationships. But it is hard to see what other concept could underlie duties to *utter* dependents. If dependence underlies those duties, and if we want a unified care ethics, then we have initial reason to see if dependence might underlie care ethicists' other claims, too.

All of this is just to show the initial plausibility of a dependence-like basis for care ethics. As indicated earlier, I will focus on my particular conception of dependence as the possible basis for the four claims of care ethics laid out in Chapter Four. This more specific argument will have to consider more deeply each of the four claims in turn. As I indicated earlier, there are two questions to ask regarding each claim. First, do care ethicists and the dependence-based principles produce the *same deliverances* regarding the claims' domains of concern? Second, do the principles provide a *good explanation* of the claims' deliverances?

### **5.3 Claim 1: Sympathetic and Contextually-orientated Deliberation**

#### **5.3.1 The Principles Call for Some Sympathy and Contextuality**

Claim 1 is:

Ethical theory should positively endorse deliberation involving sympathy and direct attendance to concrete particulars.

Do care ethicists and the dependence-based principles produce the same deliverances regarding sympathy and direct attention to concrete particulars? Recall that Claim 1 is about how we ought to reason. It does not necessarily imply anything about what makes the results of moral reasoning true. Additionally, Claim 1 is silent on exactly when and why we should engage in various kinds of moral reasoning, so we cannot ask whether the principles endorse certain kinds of deliberation in the same *cases* as Claim 1. We can, though, ask whether the endorsement of sympathy and contextually-orientated deliberation arises naturally out of the principles in a good number of cases. We can also ask whether the use of sympathy and a wholly contextual orientation is, by care ethical lights, appropriately constrained by the considerations evoked in the principles.

There is obviously a circumscribed version of Claim 1 contained within the Dependency and Coordination Principles, if we consider them in light of very plausible

assumptions about human reasoning and environmental circumstances. First, consider that, for most humans most of the time, the *best way* to fulfil the principles' duties isn't to refer to the principles, but rather to have internalised them, so that one acts, generally, in accordance with them without explicitly entertaining one's actions in that way. This will allow for faster, more efficient action.<sup>86</sup> Second, attending to such particularities will sometimes be *essential* for fulfilling the interests that are mentioned in the principles. Some important interests can be fulfilled only if the agent does not entertain some abstract principle—consider the important interest of being loved, for example.

Can these reflections apply to collectives as well as individuals? It might seem that collectives—characterised as they usually are by formal decision-making procedures—do not have scope for the kind of sympathetic and particularistic deliberation Claim 1 calls for. If they can't, then Claim 1 seems at a far remove from my account of both the Dependency Principle (when applied to collectives) and the Coordination Principle: on my account, the principles generate duties for agents on whom Claim 1 can make no demands. But surely Claim 1 applies to all agents. Thus it either applies to collectives, or I am wrong that collectives are moral agents.

Noddings, for example, seems to deny that collectives can have moral reasons on the basis of Claim 1, asserting that “[i]n a deep sense, no institution or nation can be ethical. It cannot meet the other as one-caring or as one trying to care. ... Only the individual can be truly called to ethical behaviour...” (1995, 29). By “ethical” Noddings seems to refer to a personal, particular, concrete, encounter between one human and another. If we view the collective as a mechanism independent of its members, then of course it cannot engage in such encounters. Aside from anything else, it lacks the necessary corporeality. But in Chapter Three I suggested that we not view collectives like this. Collectives are just a set of individuals arranged in a certain way. And these individuals have just the kind of particularity and concreteness that Noddings sees as necessary for ethical action, and that are necessary for sympathetic and particularistic deliberation about how best to care in ethical encounters. On Chapter Three's analysis, collective acts are a sum of individual acts that are performed within, and because of, those individuals' roles within a collective procedure. On this understanding, collectives can distribute to members the role “engage sympathetically with this individual,” and when a member acts within this role with a view to pursuing the collective's goal, then the collective has, in a sense, engaged sympathetically with that individual.

---

<sup>86</sup> Of course, this is probably true of most principles. To that extent, these reflections do not show that the Dependence and Coordination Principles are align with care ethics' judgments on sympathy and contextually-oriented deliberation *better than* other principles. This comparative question will be addressed in §5.3.3.

Take, for example, a neglected and emotionally battered child. Assume that the fulfilment of an important interest of this child requires that he receives sympathy, personalised attention, and a reluctance to catalogue his situation within some generic kind. Individuals in a collective context might be capable of using a collective decision-making procedure to distribute roles to members that are sufficient for producing these responses—for example, a local social welfare officer might have the ability to use her role in the social welfare department's decision-making procedure to distribute roles according to which another welfare officer will meet with the child once a week. This officer might in turn use his role in the procedure to distribute the role of meeting with the child to a child psychologist, and so on. This is all rather cold and bureaucratic, which is presumably what Noddings means when she says that no institution or nation can be ethical. Yet it produces the kinds of sympathy and particularism in deliberation that care ethicists endorse.<sup>87</sup>

### **5.3.2 The Principles Call for Enough Sympathy and Contextuality**

Even if the principles require *some* sympathy and contextuality, do they support sympathetic, contextual deliberation in the *same cases* as care ethicists? It might seem that care ethicists surely advocate it more often. That is difficult to assess since care ethicists tend not to say exactly when the different kinds of deliberation are required. But it is important to note that care ethicists clearly do not mean for agents to reflexively attend to others, with complete disregard for the costs to themselves or others, and without consideration of what other kinds of caring arrangements might be better, or what would happen if care was always given in this kind of situation. That would go against the widespread concern of care ethicists to restructure the social and political division of care work (particularly the gendered division of such work) (Tronto 1993; Kittay 1999). (Though early Noddings (1984) seemed to endorse an ideal of selflessness in carers, this is scaled back in her later work (2002)). If care ethicists want scope for *critiquing* general patterns of care, they will have to endorse more impartial and abstract modes of deliberation more often than some of their comments in support of Claim 1 suggest.

Of course, this critiquing might not require the hard-and-fast application of principles, either. But it does seem to require attuned sensitivity to the various *types* of

---

<sup>87</sup> We might be tempted to say, still, the collective *as such* cannot deliberate in the way care ethicists demand. Only their natural individual members can do this. Yet only natural individuals can sign contracts, and we do not take that mean that natural individuals cannot sign contracts within and because of their role within a collective, because the collective has decided that this is something the collective should do. It is unclear why we should think sympathy and contextual reasoning are restricted in a way contract signing is not.

needs and competencies that arise across cases, so that care can be given more effectively overall. Such abstraction likely circumscribes the intensity of sympathy and particularised attention that should be employed in regard to one given case, in favour of more impartial and abstract reflection on the kinds of care practices it is generally worthwhile to have across a range of cases—at least, assuming one cannot attend sympathetically to the particulars of *all possible* instances of care when deliberating about social practices of care. This will require deliberation on the basis of abstract principles, at least in quiet moments of reflection about one's (or one's society's) general practices of care—and perhaps also in the heat of the moment, when one finds oneself caring in ways that are contradictory to the general pattern of care one endorses.

For example, if a volunteer at a homeless shelter becomes so sympathetically engrossed in the details of one client's case that he is compelled to do things other volunteers could do better, or for which he would incur unreasonable costs if done in all like cases, then he should step back from sympathetic engagement and into principled deliberation. With this, care ethicists can (and should) agree. (As Tronto puts it, “[i]f the preservation of a web of relationships is the starting premise of an ethic of care, then there is little basis for critical reflection on whether those relationships are good, healthy, or worthy of preservation” (1987, 660). The principles from Part I give a good basis for such critical reflection.)

Thus though the principles will endorse such deliberation more infrequently than some care ethicists would endorse, I suggest that care ethicists should revise their endorsements accordingly. Advocating sympathetic and contextual deliberation when the principles suggest is true to care ethics' guiding concern not only to provide care for dependents, but to ensure “*that all* [including caregivers themselves, through principled critical reflection on their care practices] *will be adequately attended in relations that are sustaining*” (Kittay, 1999, 113, emphasis in original).

Yet one might press the following objection: sympathetic and contextually-oriented deliberation is important not only for those would-be carers who are best-placed to fulfil an important interest, but for those carers who are not best-placed. So the explanation for why such deliberation is important can't be the Dependency and Coordination Principles. These can only recommend that the best-placed agent engage in such deliberation. This objector agrees that non-sympathetic, non-contextually oriented deliberation will sometimes be necessary for revising our care practices, but points out that even such revision will lead us to endorse a sympathetic, contextual deliberative practice for *all* (or at least many) agents, not just those that are *best-placed*.

It is true that the black letter of the principles only make prescriptions for best-placed agents. They say nothing explicit about the duties of non-best placed agents, so seem to imply nothing about how those agents should deliberate. But when it comes to actually *applying* the principles, things will be quite different. If the Dependency and Coordination Principles are true, *all* moral agents should be closely attentive to the complex details of situations they find themselves in, in order to establish whether they or some other agent is best-placed. This will require attention to others' important interests, to what measures oneself and others have for fulfilling it, to whether those others are likely to actually discharge a duty if they have one, and so on. These inferences are largely about the psyches of other agents—and situations in which important emotional interests are at stake will require inferences about the psyche of the dependent, as well. Actually applying the principles, then, will require sympathy and deliberation from all agents who have reason to believe they might have a dependency or coordination duty in that context. They will not just require from the agent who turns out to be best-placed.

### **5.3.3 The Principles Explain Claim 1**

Do the principles provide a *good explanation* of Claim 1's deliverances? The principles' explanation is that these kinds of deliberation enable efficient, spontaneous action and the fulfilment of some interests that cannot be fulfilled from a motive of duty. They also facilitate agents' accurate assessment of whether they or someone else is best-placed to fulfil a given important interest. One might think that all kinds of principles can produce exactly these explanations: why think that the principle that the *best-placed agent should fulfil the interest* is the right one?

Simply, the idea that the best-placed agent should fulfil the interest is true to care ethicists' own justifications for eschewing principles in deliberation, which we saw in Chapter Four: namely, so that we can *better respond to persons' needs*. The injunction is that these needs should be responded to in the most dependent-centred way possible, to make the dependent's life better. The principles and care ethics agree that it is the needs, or important interests, of persons that pressure us to eschew the explicit entertainment of principles.

Rival explanations would suggest that we should engage in sympathetic and contextually-oriented deliberation because this would enable us to abide by the principles of, say, reciprocity, or contribution, or promise-keeping, rather than the dependence-based principles. These principles state roughly that one should give what one gets, or that one should fix what one breaks, or that one does what one has assured on will. Abiding by

these principles might well require sympathetic and contextual deliberation. But they are not care ethicists' concern. If the concern is *responding to needs*, then, barring evidence to the contrary, there is good reason to think that this suggests a concern for responding to needs *in the most effective way possible*. This will involve the best-placed agent (of all willing agents) responding to them.

One might object more generally that principles cannot be used to justify Claim 1. Using principles to do this seems to produce a catch-22: if an agent has the sympathetic, context-orientated deliberative practices that care ethicists advocate *because* those practices are required by the principles, she seems to have undermined those practices. Stocker (1976) claimed that this demands “schizophrenia” from moral agents. We can think of the objection this way. If care ethicists accept the principles as criteria of rightness, agents are left in a dilemma. On one horn, the agent entertains the thought that not attending to principles is demanded by the principles. By entertaining this thought, she seems to fail to meet the demand not to attend to the principles. So she fails to have motivations that the principles demand. On the second horn, if she uses sympathetic or otherwise solely contextually-orientated deliberation, without entertaining the fact that the principles demand this, then the advocate of the dependence-based principles must negatively evaluate her internal states: she’s doing the right thing, but she doesn’t know it, so she shouldn’t be praised. Either way, it seems she has done wrong. Or, perhaps on the second horn, she should be praised, but the principles are thereby self-effacing: they demand that agents not entertain them (at least, they demand this sometimes).

I suggest that we take the second horn of this dilemma. Consider what the horn really involves. If the principles are self-effacing, they imply that agents sometimes can do right only if they don’t think about the ultimate justifications. This is not to say that they should *never* think about the ultimate justifications (in moments of “quiet reflection,” for example). It is just that this is not always best. Simply, one can miss a goal by aiming at it: one can fail to ensure that dependents receive care by having that value in mind. In this way, responding appropriately to the principles’ truth, through sympathetic and solely context-orientated care, is akin to falling asleep. The best way to fall asleep is not to think about the need to sleep. A goal gets realised that is valuable, but the goal could not be realised if its value was the agent’s focus. (Elster (1985, Part II) calls these “states that are essentially by-products.”)

Some view such self-effacement as a huge cost for moral principles (e.g. Anderson 1993, chs 3–4). I suggest, though, that if the principle is explanatory—if it give us compelling answers as to why agents should act in some way, even if agents should not always entertain that “why”—then self-effacing principles are an important addition to an

ethical theory. They allow us to say more than just “care is appropriate in this case” or “here’s how to decide that care is appropriate in this case.” They can tell us *why* care is appropriate in that case.

The dependence-based principles are, in some cases, self-effacing. But not in all. I have suggested that care ethicists should endorse drawing principles into deliberation more often than some of them seem to. Yet the principles allow for much more sympathy and contextual attention than may first meet the eye, since these modes of interpretation might be important for figuring out how important various interests and values are, and thus for figuring out whether the principles’ antecedents are true. Moreover, by giving us an *explanation* for why different kinds of deliberation are required in different cases, the principles can exactly *precisify* when different kinds of deliberation are required and they can *unify* the deliberative prescriptions of care ethics in a well-motivated way.

## 5.4 Claim 2: Relationship Importance

### 5.4.1 The Principles Generate Some Relationship Duties

Claim 2 is:

To the extent that they have valuable meanings for (including effects on) individuals in the relationship, personal relationships ought to be (a) treated as moral paradigms, (b) valued, preserved, or promoted (as appropriate to the circumstance at hand) and (c) acknowledged as giving rise to weighty duties.

Call (a), (b), and (c) the “relationship duties.” Again, I need to address two questions. First, do care ethicists and the dependence-based principles produce the same deliverances regarding relationship duties? Second, do the principles provide a *good explanation* of Claim 2’s deliverances?

Let me start by explaining how the principles generate relationship duties. First, in virtue of the fact (when it is a fact) that the participants in such relationships value each other (or value the relationship), one relative’s fulfilling relationship duties has a certain kind of significance to the other relative that the same duties, discharged by someone else, would not. Our personal relatives—whether family, friends, or spouses—are thus uniquely capable of fulfilling certain interests. Since each one of our friends and family is unique, and each relationship with each of them is unique, they are likely to be best-placed to fulfil slightly different important interests of ours. (Along similar lines, Keller (2006) argues that

children can give their parents “special goods.”) Personal relatives can fulfil these interests only through believing in—and discharging—the relationship duties.

Second, there are some interests that relatives know best how to fulfil for one another or are most motivated to fulfil for one another (Jackson 1991; Pettit and Goodin 1986). Although others could in principle fulfil them, personal relatives will do so most efficiently or others are known not to be willing to fulfil them. (And we must exclude those unwilling others as rival best-placed agents when considering whether personal relatives are best-placed, according to the principles.)

Third, discharging the relationship duties involves fulfilling many of the *agent's* important interests that could not otherwise be fulfilled: the joy and satisfaction of knowing the relative deeply, of seeing him or her nurtured by one's own hand, and so on. This will decrease the cost to the agent of taking such measures—and so will increase the value of the measures overall, regarding dependent and agent.

Let us assume that—for many people at least—some of the interests that can be best fulfilled through relationship duties are towards the “important” end of the important–unimportant interest continuum. Then, according to the principles, relatives should discharge the relationships duties in order to fulfil those interests, either acting alone (under the Dependency Principle) or working together to do so (under the Coordination Principle).<sup>88</sup> These dependency and coordination duties arise in virtue of the relationship, insofar as it is the relationship that renders the agents best-placed. And these dependency and coordination duties are discharged *through* discharging the relationship duties. The relationship duties are then an upshot of dependency or coordination duties.

Notice as well that these duties can be held by collectives. Suppose a group of friends organises themselves to go hiking once a month. They have a particular procedure (perhaps discussion-based consensus) that is employed each month for deciding where to hike, for how long, how to get there, and so on. The collective has certain moral obligations to its members to discharge relationship duties to those members. This is because *the collective* can fulfil certain important interests of each member, such as the interest in a sense of belonging, in having a sense of shared history, and in experiencing the shifting tides of group dynamics. While each individual in the collective might have similar interests fulfilled in each of the dyadic relationships they share with each other member, there are additional interests that can be fulfilled only in a *group* context, by a well-oiled

---

<sup>88</sup> Often, two or more adults (such as parents, an extended family, a group of friends, or a group of community members) meet the Coordination Principle for fulfilling certain interests of someone, such as the interest in a stable home life or a sense of belonging. These agents then incur coordination duties, under the Coordination Principle, to *work together* for the fulfilment of that interest.

*collective* of friends. And these interests may well require that the collective treats each member (in his or her capacity as an individual, rather than in his or her capacity as a member of the collective) in the ways demanded by the relationship duties. These duties are not necessarily made explicit in the collective's decision-making procedure. They rather arise involuntarily, out of the individuals' dependence on the collective. This is an example of a relationship between an individual (qua individual, not qua member) and a collective (to which she belongs) generating relationship duties for the collective.

Another reason for the Dependency and Coordination Principles' production of relationship duties is the effects that taking personal relationships as moral paradigms is likely to have on non-personal dependence relations. In personal relationships, the fulfilment of dependence-based duties usually brings with it a certain joy and fulfilment for the agent. This is not always so in non-personal relationships. By taking the same kind of attitude (if not to the same extent) to duties in non-personal relationships that we do to duties in personal relationships, we can make care in non-personal relationships less costly for ourselves. This can help us to fulfil our dependence-based duties in non-personal relationships.

#### **5.4.2 The Principles Generate Enough Relationship Duties**

Of course, discharging duties within personal relationships is not always so heart-warming. Not all family members, say, value each other or the relationship; not all are able to give one another unique goods; not all are particularly motivated to care about and for each other. If these people discharged the relationship duties for each other, perhaps no important interests would be fulfilled. And they certainly do not seem *best-placed* to fulfil those important interests for each other. The Dependency and Coordination Principles give rise to duties—including relationship duties—only when discharging relationship duties is the best way for agents to fulfil important interests that they are best-placed to fulfil.

Thus it might seem that a care ethicist will see relationship duties as arising more often than the principles allow. This is the same problem as that which we encountered with Claim 1: it seems that the claim of care ethics is going to produce prescriptions even when the antecedent of the Dependency or Coordination Principle is not satisfied. However, I suggest that the best version of care ethics should similarly recognise relationship duties if and only if there are some important interests that one relative is best-placed to provide for the other—important interests being, in this context, important positive effects on, meanings for, or value to, individuals. This is not something all care

ethicists will readily assent to. But, I suggest, it is something that they *should* assent to, given their underlying concerns and the aim (presumed in this chapter) of having a precise and unified theory.

This might seem objectionably dismissive of personal relationships where participants happen not to be *best-placed* to fulfil other participants' important interests. Yet it will produce more relationship duties than one might think. Recall from Chapter Two that we should understand “best-placed” diachronically. The ability to give our relatives unique goods, to know best how to give them non-unique goods, or to be most motivated to do so, must be considered in this diachronic way. Thus, while it might be the case that *right now* a given parent-child (or parents-child, or family-child) relationship does not have important positive effects on, meanings for, or value to, individuals—so the Dependency and Coordination Principles seemingly do not generate relationship duties for its participants—it might still be true that the parent (or parents, or family) are best-placed, diachronically, to fulfil the interest, because of the expected value of them making themselves such that they are best-placed (synchronously) to fulfil the interest. In such a case, the Dependency and Coordination Principles would hold in a diachronic way. This may generate duties to reconcile estrangements, to re-ignite lost enjoyment, to keep in touch over distances, and so on. If there is absolutely no important interest—even if that interest is very specific—that a relative is best-placed to provide for us—either diachronically and synchronically—then we should be reluctant to say there are relationship duties in that relationship.

There is another way in which the Dependency and Coordination Principles might lead to more partiality to personal relatives than one might think. Perhaps, given plausible assumptions about human psychology, it would be prohibitively costly for us not to use personal relationships as paradigmatic moral interactions; to assign personal relationships added value; and to see our general duties to everyone as being more weighty in the case of personal relatives. Perhaps we just cannot be motivated to care to a comparable degree about or for non-relatives; perhaps personal relationships cultivate in us compassionate dispositions that are arguably necessary for us to care at all about or for non-relatives; perhaps, without recognising relationship importance, we would be depressed and unable to provide any care to anyone. It might just be psychologically or socially impossible for us not to care, primarily, about and for personal relatives.<sup>89</sup> If that's true, and if ought implies

---

<sup>89</sup> The fact that we are inescapably embedded in social life, and that we cannot abstract away from that embeddedness when it comes to moral practice, plays a core role in feminist and care ethical critiques of “traditional” moral theories. See especially Young 1990, esp. pp. 104–5; Benhabib 1987; Gilligan 1982.

can, then the commands of the Dependency and Coordination Principles will be constrained to allow such partiality.

Again, this plausible assumption about human psychology is only a general rule. And again, it has the following implication: if, for a given agent, it would not be prohibitively costly to ignore relationship duties; and if the fulfilment of other of that agent's dependence-based duties hinge on the agent's ignoring relationship duties; then the agent has a duty to ignore relationship duties. Which is just to say: if no very important interests hinge on your treating your personal relationships as important; and some very important interests hinge on you *not* treating your relatives as important; then you shouldn't treat your relatives as important. With this, I think, care ethicists should, and would, concur.

One might think this cannot possibly go far enough. Surely the principles will, much of the time, require that we not focus on personal relationships, but rather dispense with personal relationships and spend the saved time at homeless shelters or doing paid work to give the pay to international development organisations. This is especially a problem for the Dependency Principle, whose third consequent states that the duties are stronger—the agent does more moral wrong if he defaults on them—if the targeted interest is more important and if the value of the agent's measures are higher regarding agent and dependent. It will often be the case that distant strangers have vitally important unfulfilled interests, where our measures for fulfilling them will realise far more value than our measures for fulfilling interests of our personal relatives. (Among numerous others, Wolf (1982, 428) makes something like this critique of utilitarianism; Singer (1972) embraces its conclusions.) That is, though the Dependency and Coordination Principles explain *some* relationship importance, they do not give us as much of it as care ethicists want.

There are two responses to this. First, duties to distant others are not anathema to care ethicists. As was discussed in Chapter Four, much recent work has been on “globalising” the theory (e.g., Held 2006; Engster 2007; Kittay 2008; S.C. Miller 2010; Robinson 2011; Tronto 2010). I suggest that care ethicists can only legitimately make this “globalising” move after they have acknowledged that something like dependence is the core underlying their theory. If care ethics is to be more than whatever claims are endorsed by the people who happen to be known as “care ethicists”—and this chapter aims to see it as more than that—then there must be an internal logic to the inclusion of *global care concerns*

---

Though see Estlund (2011) for an argument for the conclusion that the fact that we are psychologically incapable of X-ing does not mean that it is not the case that we ought to X.

within the tradition. The Dependency and Coordination Principles give a coherent story for how that extension of the theory is possible. Without the general structure of these principles as part of the theory, it is not clear what can justify this globalising move from *within* the theory. If this demands that we do more for distant others than many of us are currently comfortable with, so be it.

The second response is to point out that there could well be distinct kinds of value in deep personal relationships, which, if forgone in favour of assisting others, would actually cause that assistance to have greatly *negative* expected value for the agent. This would render the assistance too costly, and require that the agent stay in the personal relationship. And these kinds of value might be wholly independent of the *important interests* those relationships fulfil. If so, those sources of value would serve as costs to the agent, which constrain the demands non-personal relationships could make on us, by the lights of the Dependency and Coordination Principles.

#### **5.4.3 Extension: Duties to Form Personal Relationships**

Having meaningful, deeply valued personal relationships is plausibly an important interest of most individuals. And the important interests of individuals are what the Dependency and Coordination Principles are concerned with. To the extent that personal relationships imbue our lives with meaning and value, we can generalise and assign them relationship importance, that is, status as morally paradigmatic, valuable, or generating of weighty duties. But it is not necessarily the *personalness* per se that gives these relationships this status, according to the principles. If the principles are right about this, then do we sometimes have duties to *take steps to form* personal relationships, in those cases where we are best-placed to fulfil another person's important interests in experiencing the meaning, value, love, and intimacy that personal relationships bring? I will argue that the answer is “yes” by dependence-based lights, and that care ethicists ought, by their own lights, to embrace this.<sup>90</sup>

Consider the most straightforward example, friendship. This example is most straightforward because friendship is not as exclusive as, say, romantic relationships. Thus developing a friendship with one person will not so often require foreclosing the possibility of this kind of relationship with others. So the costs of developing these relationships will often not be so high. (Though similar considerations apply to romantic partners.) Consider the current distribution of valuable, meaningful friendships. Some people have more

---

<sup>90</sup> I provide a similar argument for duties to form new personal relationships in Collins forthcoming.

friends than others, simply because they are more witty, or more spontaneous, or more intelligent. Of course, the value of friendship largely lies in the enjoyment of characteristics like funniness, spontaneity, and intelligence. So it would be counterproductive to say we should befriend those who we find overly serious, predictable, and dense, because we would be unable to fulfil the interest these people have in (the right kind of) friendship. Nonetheless, there is room for moral imperatives to try to change the relationships within which we produce the unique goods that friendship brings.

Specifically, by the lights of the Dependency and Coordination Principles, the following is true: we have a duty to adopt actions and attitudes with a view to getting into a personal relationship with a person, if we are best-placed to get into a personal relationship with that person that would fulfil the important interest they have in such a relationship. This is just a special instance of the diachronic fulfilment of dependency and coordination duties. It applies, for example, to the mother who doesn't (but could grow to) love her child in ways that no one else can. Here it applies through the Dependency Principle. It also applies to the group of friends who, through mutual responsiveness, could draw a newcomer into their circle—where the action of any one friend would, we assume, not be enough to bring about the unique kinds of goods that come from being part of a *group* of friends. Here it applies through the Coordination Principle.

There are a few caveats to place on this point. First, this duty is just to take steps with a view to getting into a personal relationship. The duty is not to *get into* and then *stay in* relationships under duress, precisely because such relationships would not fulfil our interests in genuine love, camaraderie, intimacy, or whatever is the particular important interest at hand. This is also because personal relationships are a two-way street, so one cannot form or stay in a relationship on one's own. Rather, one can only take actions that make it *likely*. Second, because we are always already embedded in personal relationships, it might be that the Dependency and Coordination Principles demand that we stay in at least some of those relationships. This is simply because of the costs—to ourselves and to our current relatives—of ending them. Third, the principles might be indeterminate about precisely with whom one should increase one's chance of getting into a relationship. So there might be scope for discretion in discharging the relationships duty. Fourth, the principles give sufficient, not necessary, reasons to take steps to form personal relationships. They do not *exclude* getting into personal relationships for non-moral reasons, as long as those reasons are constrained by moral principles where appropriate.

Care ethicists, I suggest, should embrace duties to take steps to form new personal relationships. That the principles have this implication goes to the heart of why

relationships matter to us: they are constituted by, constitute, and produce goods that uniquely fulfil a wide range of interests, for which we are then dependent on our relatives.

One might object that this implication of Part I's principles demands that we initiate personal relationships for the wrong reasons. These reasons are wrong because they discourage, or even preclude, *real* friendship. Real friendships, and real friends, must be valued not because of their relationship to a moral duty. (Tronto 1993, 105; Held 2006, 80; Stocker 1976, 462; and Blum 1980, 142–3 each make related points about the case of friendship.) Here again it is important to separate justification from motivation, as we did when discussing Claim 1: the principles may well demand that these relationship-forming steps not be entertained as a duty. Additionally, in many cases, moral reasons are perfectly good motivations for forming friendships. If, months after moving to a new city, a friend tells you “you know, I made an effort to be especially considerate of you when you first moved here because you seemed lonely,” this would not seem like the wrong reason. On the contrary, it would seem a paradigmatic case of sympathy, contextual sensitivity, caring about, and a recognition of friendship's importance—all things on which care ethicists rightly place great value.

#### **5.4.4 The Principles Explain Claim 2**

There remains the question of whether dependence gives a good explanation—a good justificatory ground—of Claim 2's truth. As initial support for dependence as an explanation of Claim 2, notice that at least one prominent care ethicists seems to presuppose its applicability. Engster argues that in a relationship between A and B, the relationship gives rise to responsibilities because either: (a) A is able to care for B at short notice and with the least costly movement of resources; (b) A knows specifically what B needs; (c) A has B's acceptance and confidence as a carer; or (d) A has an affection for B that makes her more motivated to care for B. Engster's suggestion seems to be that the more of these attributes a relationship has, and the more strongly that attribute is instantiated, the stronger a responsibility A has toward B. On the basis of (a)–(d), Engster asserts that our caring responsibilities are foremost to ourselves, then to family and friends, and then to community members and compatriots (2007, 54–7). “Distant strangers” come fourth and last because we are “least well able and least well positioned to deliver good care to distant strangers” (2007, 57). This hierarchy justifies Engster's “principle of subsidiarity,” which states that “we should shift the actual delivery of care whenever possible to the most local and personal levels. We should care for others whenever possible by enabling them to care for themselves” (2007, 58). The principle of

subsidiarity is based on the fact that care can be most efficiently and effectively administered at the local and personal levels.

More important than one hint in the literature, though, is dependence's intuitive explanatory force. In §4.4.2, I considered some rival explanations as presented by Kolodny and Scheffler. These based relationship duties in the relationship's having the right kind of history, or in its being non-instrumentally valuable. But these did not tell us why some histories (e.g. between friends) are more important (paradigmatic, valuable, normatively weighty) than others (e.g. between dentist and patient), or what a good reason is to non-instrumentally value a relationship. They seem to appeal to some mysterious power that personal relationships have in a brute way, even if they are bad for those who are in them.

By contrast, a dependence-based explanation just asks us to recognise the vitally important goods—support, sympathy, attention, compassion—that we receive from those who love us. It asks us to consider that that we *depend* on them for these goods, that we are hugely vulnerable to their excluding us from them, that we would have no one else (or no one else nearly as worthwhile) to turn to if they left us. The formal and calculating nature of the Dependency and Coordination Principles looks out of place here, and rightly so: we should not entertain these principles every time we consider comforting a friend, phoning a parent, or embracing a lover. But the comfort, phone call, and embrace are ways of fulfilling important interests of these people—important interests that, often, we are best-placed to meet (perhaps in mutual responsiveness with friends, siblings, or so on; or as part of a formal collective such as a club or sports team).

Of course, sometimes strong social norms inhibit us from altering relationships in the ways the principles demand, leaving us stuck, for example, with a web of family relationships that is in the interests of no one. The Dependency and Coordination Principles do not—and, I suggest, the best version of care ethics *should* not—endorse such relationships, if it is really the case that they do not fulfil important interests. It would be surprising, however, if such relationships did not make their participants best-placed to fulfil at least some of their relatives' fairly important interests. In connection to our family relationships, consider, for example, our interests in receiving respect from other members of society, having a sense of home and belonging, or having a socially-developed conception of the good life. Such family relationships are likely to fulfil at least some of these interests. And for many of them, our family will be best-placed to fulfil them.

## 5.5 Claim 3: Caring About

### 5.5.1 *The Principles Call for Caring About*

Claim 3 is:

Care ethics sometimes calls for agents to have attitudes that: (i) have as their object something that has (or has the prospect of) a decent life, or something that might affect something that has (or has the prospect of) a decent life; and that (ii) are a positive response (e.g. promoting, respecting, revering) to that life (or prospect); and that (iii) lead the agent's affects, desires, decisions, attention, or so on to be influenced by how the agent believes things are going with the life-bearer (or prospect-bearer).

Claim 3 does not say in what circumstances agents have these duties. We thus cannot decidedly assess whether care ethics and the dependence-based principles produce the same deliverances regarding caring attitudes. We can, though, assess whether and when the Dependency and Coordination Principles would generate a duty to have caring attitudes, and assess this in light of educated guesses about the circumstances within which care ethicists would posit such duties.

First, does the Dependency Principle generate duties for agents to have such attitudes? Yes, if and only if the following conditions hold: the attitude is the most efficacious means the agent has for fulfilling some person's important interest; if adopted, the attitude would have a proportionately high likelihood of fulfilling the interest (i.e. proportionate to the interest's importance); adopting the attitude would realise positive iterative expected value regarding agent and dependent; a policy of adopting like attitudes in like cases would realise positive aggregate expected moral value regarding agent and dependents; and the agent's adopting the attitude would realise no less expected value regarding agent and dependent than another agent's taking measures to fulfil the relevant interest.

What about the Coordination Principle? It demands caring attitudes if and only if: they are part of the most efficacious set of steps an agent could take in responsiveness to others to realise some state of affairs in which important interests are fulfilled; where adopting the attitude would create positive iterative expected value regarding agent and dependent; where a policy of adopting like attitudes in like cases would create positive aggregate expected moral value regarding agent and dependents; and finally where no other

set of agents will take steps (towards the state affairs) that would produce higher expected value regarding agent and dependent.

These conditions will hold more often than one might think. To see this, recall Chapter Four's discussion of the moral value of emotions. Emotions can be at least instrumentally, and plausibly also non-instrumentally, valuable. In their instrumental capacity, they can be part of the *most efficacious* set of responsive measures taken to fulfil someone's important interests, even if they are not strictly *necessary* for that fulfilment. Adopting attitudes might reduce the cognitive and motivational effort, and thus the cost to the agent, of taking interest-fulfilling measures. This would increase the measures' expected value for the agent. I said in Chapter Two that agents have to consider the cost they will incur diachronically. If cultivating certain attitudes will make the total diachronic value positive, then part of fulfilling the dependency or coordination duty will be cultivating those attitudes. Additionally, emotions plausibly have non-instrumental value for both agents and dependents in some cases: it is better for aged Lesley to be cared for by loving Barbara than by compassionless Graham.

### **5.5.2 The Principles Call for Enough Caring About**

Once again, however, it may seem care ethicists are concerned to prescribe a much broader range of attitudes than can be required by the Dependency and Coordination Principles. According to the principles, the interest must be *important* and the agent (or set of agents) must be *best-placed* to fulfil the interest. Ruddick (2002, 222) critiques Kittay's dependency-based theory of care along similar lines: a caring attitude might be demanded just when my friend is upset, even I am not best-placed to help her (say, because she has many friends who know her better) and even if her interest is not particularly important. Along the same lines, Held insists that "when we ... understand how increasing levels of affection, mutual concern, and emotional satisfaction are valuable, we can aim at promoting care far beyond the levels of necessity" (Held 2004, 63). To put an example to this thought, perhaps I should feel (and express) sadness for my colleague when I find out her mother has died—even though I am not best-placed to make her feel better through empathy (maybe her spouse is better-placed).

If the Dependency or Coordination Principle does not generate a duty for me to feel and express sadness for my colleague, then this is because the importance of her interest in my sadness is not sufficiently important compared to the costs to me—such as having discretion over when to feel sad on another's behalf, or caring more about my own problems than hers, or so on. That is, condition (2) or condition (3) of the Dependency

Principle are not satisfied: her interest's importance is not proportionate to my capacity to fulfil it; or my taking measures to fulfil it would realise negative value with regard to myself. So I do not have a *duty* to empathise.

Yet if care ethicists merely claim there is *some* positive reason to feel sad for my colleague, or that it would be *admirable* (but not a *duty*), the principles might agree, depending on the details of the case and what constitutes the costs and benefits. A proponent of the principles might declare it *good* or *admirable* or *decent* for me to be sad for my colleague, by doing so at my discretion or in addition to caring about my own trivial problems. They simply deny that these values weigh up to give me an *duty* to do so—that is, a *weighty* moral reason that requires *strong* defeaters in order to block an all-things-considered obligation to do that thing. The principles only deliver verdicts on duties: they are silent on weaker kinds of reasons. And it is certainly in keeping with their spirit—concerned as they ultimately are with the effective fulfilment of important interests—to acknowledge weak reasons to fulfil interests in cases where, say, the cost to the agent is high but the interest is fairly important.

Alternatively, if care ethicists want to say that there is an moral *duty* to feel sad for my colleague, while the principles deny this, then care ethicists must be latching onto something other than the importance of her interest in my empathy, and my being best-placed to fulfil that interest, as the reason for the duty. What might this be? Take the perspective of the colleague. Why value the empathy? Plausibly, we appreciate others' empathising with our personal tragedies because this reminds us that we are not alone, helps us to see the bigger picture, and suggests that others have suffered similarly and have managed to overcome the worst of the grief. All these look like very important interests.

So if care ethicists recognise an all-things-considered duty while the principles do not, then it must be because care ethicists deny that my duty to feel sad for my colleague derives from my being *best-placed* to fulfil my colleague's important interest. Care ethics and the Dependency Principle do come apart in this way: if I am not best-placed to fulfil my colleague's interest on my own, the Dependency Principle cannot say there is a duty. Yet care ethicists might say that I do have a duty—and perhaps rightly so.

However, the Dependency Principle is not the only dependence-based principle. I suggest that we view this kind of case through the lens of the Coordination Principle. Presumably, every little bit of empathy from every person helps, such that my duty to fulfil my colleague's interests becomes part of a coordination duty—I and others together can help her to overcome her grief, and I have a duty to act responsively to others with a view to achieving that end. I am not best-placed on my own to fulfil her interest in empathy—after all, the final interests that my empathy is instrumental to are quite significant

realisations regarding the role of grief in my colleague's life and in the world. I cannot bring about these realisations in my colleague on my own, let alone am I best-placed to do so. But I and others combined are best-placed to do so, if we act responsively in the right way. So my duty to be empathetic to my colleague can be straightforwardly construed as a coordination duty.<sup>91</sup>

We might instead imagine that the department as a whole—as a *collective*, rather than a mere set of responsive individuals—has a duty to empathise with, or otherwise care about, the colleague's bereavement. That is, the department has a dependency duty. Or, perhaps more plausibly, the department as a collective is one of several agents—including the colleague's friends and family members—who each have a coordination duty regarding her interest in empathy. Either way, the department's *collective* empathy duty would amount to a duty for individuals to use their role within the department to use its decision-making procedure to distribute roles sufficient for expressing the department's empathy. For example, the head of department might send an email to staff explaining the situation, thus using his role to induce the empathetic sadness in members. The department itself might also have a duty to have and express this attitude. This duty would entail members each having a duty to act within their role with a view to expressing the department's empathetic sadness—perhaps by organising for everyone in department to sign a condolence card, for example.

Even given Ruddick and Held's suggestion, then, I suggest that the conditions of dependence—captured in the antecedents of the Dependency Principle and, crucially, the Coordination Principle—provide a good specification of the “sometimes” in Claim 3.

### 5.5.3 The Principles Explain Claim 3

Finally, we must consider whether dependence gives a good explanation—a good justificatory ground—of Claim 3's truth. Consider what is entailed by a duty to care about, as I analysed the concept in Chapter Four. The kind of caring about that care ethics calls for is one that has as its object something that has (the prospect of) a decent life, where the caring is a positive response to that life (or prospect). Why should it matter that we care about such objects?

---

<sup>91</sup> We can imagine this being the case even if I do not know who I am being responsive to. Maybe the bereaved person is not my colleague but a distant friend. Suppose I do not know anyone else who knows her, but I know I am not best-placed to help. Nonetheless, I might act responsively to the other empathisers insofar as my beliefs about their reasons to empathise with her affect the way I empathise with her, in a way that is a positive response to their reasons to empathise (e.g., I do not try to thwart their efforts).

Care ethicists might think that something other than dependence grounds these duties—for example, one’s having an associative relationship with the object of the attitude, or having voluntarily taken on the responsibility to have these attitudes, or having caused the need for the attitude, or expecting that others will reciprocate the attitude, or similar. (They might instead think that different duty bases apply to different contexts, but in that case they forgo having a unified explanation. A unified theory is explicitly what this chapter is trying to provide.)

Any one of these rival accounts would fail to vindicate certain strong intuitions about duties to care about, where those intuitions are central to care ethics. Any unified account of care ethics must attend to care ethicists’ emphasis on the involuntariness of, and lack of reciprocity regarding, many of our attitudinal duties. And any unified account must vindicate duties to utter dependents.

Unifying care ethics on the basis of voluntary assumption, or contribution to harm, or reciprocity, will not achieve this. Indeed, care ethicists have argued that this is an irresolvable problem with these sources of duties (Bhandary 2010; Kittay 1996; Noddings 2002, ch. 4). Similar concerns apply to association: care ethics can use association as the unified justification of duties to care about only by denying duties to care about those who are not associated with any moral agent in the right way. But care ethicists generally endorse agents’ expanding their spheres of concern to include those with whom they are not associated (Held 2006; Ruddick 1980, 123; Tronto 1993). This concern relates to that raised in relation to Claim 2: to have association as the unified explanation of duties to care about, care ethicists must forgo the possibility of duties to care about persecuted or terrorised people with whom one is not associated. Needless to say, dependence does not suffer these pitfalls.

Additionally, care ethics is usually *contrasted* with justice, where justice includes such duty bases as explicit contracts, voluntary assumption, reciprocity, contributions to harm, and so on (Kittay 1996, 232; Kittay 1999; Held 2005; Held 2006, 15–17; Robinson 1999, 23ff). If these duty bases are characterised as part of *justice* rather than as part of *care*, then duties with these bases are not best construed as care ethical.<sup>92</sup> After we subtract the sources of duties that care ethicists generally associate with justice theorists, the sources of duties left over seem to be association and dependence. In discussing Claim 2, I gave some reasons for thinking that paradigmatic sources of “associative” duties—personal relationships—are best construed (from within the framework of care ethics, at least) as

---

<sup>92</sup> Note that the Dependency and Coordination Principles are not themselves principles of *justice*, as care ethicists typically evoke that category. Justice should here be understood as grounded in the value of reciprocity and liberty—two values that are not relied upon by the Dependency and Coordination Principles.

sources of *dependence-based* duties. If these arguments were sound—and if we want a care ethics that can include and explain obligations to those we do not know personally—then we should rule out association as the one unifying explanation of duties to care about. We seem left with dependence as a plausible unified explanation for duties to care about. And given that dependence gives us a good number of variously weighted moral reasons to have attitudes, it is a good candidate to play this role.

## 5.6 Claim 4: Caring For

### 5.6.1 *The Principles Call for Caring Actions*

Claim 4 is:

Care ethics sometimes calls for agents to perform actions that (i) are performed under the (perhaps tacit) intention of fulfilling (or going some way to fulfilling) interest/s that the agent perceives some moral person (the recipient) to have; (ii) where the strength of the demand is a complex function of the value of the intention, the likelihood that the action will fulfil the interest, and the extent to which the interest is appropriately described as a “need.”

Again, we cannot assess whether the Dependency and Coordination Principles will generate duties to perform caring actions in *all and only* cases where Claim 4 does, as Claim 4 does not specify exactly when these exist. At a minimum, though, in order for the deliverances of the principles and of Claim 4 to coincide, it must be that the Dependency and Coordination Principles generate duties to perform *caring* actions, as I have analysed that concept. It might seem that the principles do not give us duties to perform caring actions, because caring actions have a certain intention, namely, fulfilling someone’s perceived interests.

Let us take the Dependency Principle first. The Dependency Principle’s consequent says nothing about intending to fulfil an interest. It states only that agents must intentionally take measures that are likely to have a certain result. But that the principle does not give us *only* duties to care does not mean that it fails to give us duties to care. In many cases—e.g., the group of surf lifesavers at the beach—part of the most efficacious measures will include something like “entertaining these measures as being facilitative of fulfilling someone’s important interest.” That the measures fulfil an important interest will often—though not always—be essential for an agent’s picking out the measures. So

often—though not always—the interest-fulfilling intention will be required, and the Dependency Principle will therefore produce duties to care.

This is true even more often in the case of the Coordination Principle. The duties that the principle produces are precisely duties to “take responsive steps with a view to p” or to “take responsive steps towards there being a collective-that-can-produce-p,” where “p” is “a non-actual state-of-affairs in which important interest(s) is (are) fulfilled.” Agents must intend to play a part in the production of interest-fulfilment. The Coordination Principle calls for responsiveness to other agents, which will require a good amount of on-the-spot guess-work and re-assessment. Again, as a matter of logical necessity, this does not *require* that the agents perceive “p” as an interest-fulfilling state of affairs. So the Coordination Principle does not generate duties over *only* caring actions. But it generates them whenever it would be valuable, or would make the interest fulfilment more efficacious, for the agent to perceive “p” as interest-fulfilling. The Dependency and Coordination Principles, then, call for caring actions.

### ***5.6.2 The Principles Call for Enough Caring Actions***

But do the principles call for *enough* caring actions? Do they call them in the right kinds of cases, by care ethical lights? There is good evidence to think that care ethicists will assert a duty to perform caring actions *at least* in cases where the Dependency or Coordination Principle holds. Kittay argues for a theory that bases responsibilities on “our *moral power* to respond to other in need” (1999, 113, emphasis added); Tronto suggests that we “assume responsibility because we recognize a need for caring, and there is *no other way* that the needs will be met except by our meeting it” (1993, 132, emphasis added); and Engster says that care for distant strangers can be de-prioritised because we are “least *well able* and least *well positioned* to deliver good care to distant strangers” (2007, 57, emphasis added). These italicised clauses do not exactly match the terminology of the principles’ antecedents, but they reflect a concern with having a high capacity to assist—perhaps more capably and at lower cost than other agents—in general.

And as with Claim 3, the distinction between duties, on the one hand, and mere moral reasons, on the other, allows us to get past the worry that care ethicists perceive mere moral reasons in a wider range of circumstances than those covered by the Dependency and Coordination Principles’ antecedents. If care ethicists say that we have merely it would be good or admirable to perform caring actions in some circumstances not covered by the principles, the principles can agree that there would be benefits from that measure being taken, while restricting their verdicts to duties. All-things-considered duties

will not be produced if those benefits are outweighed by costs to the agent, costs to the dependent, or strong defeaters.

And if care ethicists assert that there are duties to perform caring actions, when the principles do not, then this must be because they are latching onto a duty basis that is a rival to dependence (such as voluntary assumption, contribution, reciprocity, or association). If they're latching onto such a duty basis, then either their theory is fragmented (i.e., there is nothing uniting the basis of the duties it produces), or they must hold the duty basis they latch onto in that particular case is what grounds *all* duties within care ethics—including those to utter dependents and the global poor. In any case, if they latch onto such duty bases (either in one particular case or across all cases), then care ethics is either no longer clearly distinct from theories of justice, or cannot obviously generate duties to strangers. (Where justice is understood as valuing liberty and reciprocity—notions rather at odds to the interest-fulfilling aims of the Dependency and Coordination Principles.)

### **5.6.3 The Principles Explain Claim 4**

Dependence gives a good explanation—a good justificatory ground—of Claim 4's truth. To see this, consider that the second component of Claim 4 is that the strength of the demand to perform caring actions is a complex function of the value of the intention, the likelihood that the action will fulfil the interest, and the extent to which the interest is appropriately described as a “need.” This suggests that these considerations are part of the basis of the duty.

These considerations are quite straightforwardly reflected in the Dependency and Coordination Principles. When the principles consider the expected value of a proposed measure (or set of responsive steps), the principles include the value of the action's intention for the agent and dependent, and the likelihood that the interest will remain unfulfilled despite the action being taken. The principles include more considerations than these, but the fact that they include them allows us to see how the principles agree with the rationales given in Claim 4 for *why* agents have duties to perform caring actions.

Additionally, working with a notion of “important interest” on which interests sit on a scale of importance allows us to say that the interests appropriately characterised as “needs” are more important than those that are not. The language of “needs” within care ethics is reflected in the principles' concern with important interests.

## 5.7 Conclusion

I started this chapter with the aim of using the Dependency and Coordination Principles to unify, precisify, and explain the four claims of care ethics, which I developed in Chapter Four. I also aimed to use care ethics to demonstrate some surprising implications of the Dependency and Coordination Principles, which might not be obvious upon first glance.

We have seen that if care ethics is to be *unified* under the rubric of dependence, this will require giving up certain possible care ethical views—such as the view that friendships or family relationships give rise to moral duties even when participants are not best-placed to fulfil some interest of other participants, or that there are moral duties to have caring attitudes even when no important interest is at stake. I hope to have shown that these are not such terrible sacrifices for care ethicists to make at the altar of unity.

We have seen that for care ethics to be *precisified* under the rubric of dependence, care ethics has some interesting upshots. For example, under a dependence framework the claims about the importance of personal relationships generate (under the right conditions) duties to take steps to *form* personal relationships that will fulfil important interests. And given how my dependence framework operates in collective contexts, we have seen that care ethicists' claims can generate duties for collectives—as well as individuals—to deliberate via sympathy and to have, and express, caring attitudes (by using the decision-making procedure to induce such attitudes in members).

We have seen that for care ethics to be *explained* under the rubric of dependence, we must allow that the issue of motivation—roughly, the reasons we consciously entertain when performing actions and adopting attitudes—should be separated from the issue of justification—the reasons why it is morally right that we should perform those actions or adopt those attitudes. Yet the justification remains in the background, and should be brought to bear on our deliberations in moments of sober reflection on our general practices. Additionally, we have seen that this explanation accords well with care ethicists' prolific use of the language of “dependence,” “vulnerability,” and “needs.”

Finally, regarding the surprising implications of the Dependency and Coordination Principles, we have seen that these technical and abstract-seeming principles make a great deal of room for reasoning by sympathy and attention to particulars, for the importance of personal relationships, and for duties to have certain emotions and attitudes.

The discussion in Part II has largely concerned interpersonal ethics—the ethics of individuals or small groups, as most of us encounter it in our day-to-day lives. Of course, we have had some recourse to briefly consider international ethics. We have seen that many care ethicists endorse a global version of their theory, and that many of them are concerned

with the social and political inequalities that arise out of certain distributions of care work. Still, the focus in Part II has been on interpersonal ethical encounters; on the implications of dependence for individuals and small groups. But the reach of morality extends further than this—and so does the reach of dependence. In Part III, I turn to consider some work that dependence can do regarding the duties of large-scale groups in international ethics.

**PART III – INTERNATIONAL ETHICS:  
THE CASE OF RESPONSIBILITY TO PROTECT**

# Chapter Six: **The Agency of States and Intergovernmental Organisations**

## **6.1 Introduction**

In Part III, I will put my theory of dependency duties and coordination duties to work in the international realm. The ultimate purpose here is to use my account to provide a unified explanatory ground for the R2P doctrine—but R2P will be introduced only in Chapter Seven. Before I can turn to it, there is a vast gulf to be crossed: from the interpersonal and micro-collective contexts I have already discussed, to the macro-collective contexts of international politics. This short chapter is designed to bridge that gulf. In Chapter Three, I explained how both dependency and coordination duties arise in relatively small-scale group contexts. In this chapter, I show how this analysis can apply to some important super-sized real-world collectives: states and inter-governmental organisations.

Of course, these are not the only large-scale transnational collectives whose moral agency we might be interested in. (International non-governmental organisations and multinational corporations are other obvious candidates.) But these are two of the most important kinds of collective in current international political practice. In any case, they are the kinds of agents that are most often referred to by proponents of R2P. The proposition that they have moral agency is therefore an important unspoken premise in the next two chapters' arguments regarding R2P. This chapter, then, aims to briefly give an account of these collectives' agency. This will prepare us for the subsequent in-depth discussion of R2P.

## **6.2 States**

### ***6.2.1 States as Agents***

I will understand states in much the way described by The Montevideo Convention on the Right and Duties of States (1933). The Convention describes states as having: “(a) a permanent population; (b) a defined territory; (c) government; and (d) capacity to enter into relations with other states.”<sup>93</sup> We might add to this that the state is the entity that has the

---

<sup>93</sup> According to Malcolm Nathan Shaw (2003, 178) and Nandasiri Jasentuliyana (1995, 20), these criteria remain generally accepted.

final authority over the enforcement of legal rights and responsibilities within a geographic territory. Adding this condition allows us to distinguish states from, say, communities living within states that own a defined territory, have a permanent population, have a government, and enter into relations with other communities of that kind.<sup>94</sup>

States meet the criteria for collective agency I gave in Chapter Three.<sup>95</sup> They have complex decision-making procedures, which produce a range of goals and a distribution of individual tasks for achieving those goals. Many people self-consciously abide by the procedure's results; have an expectation, commonly known to be held by each member, that enough other role-bearers will abide by the procedure's results; and are able to receive sufficient information about the procedure's outputs for them to abide by it. The collective decision-making procedure is separate from, and additional to, the sum of those of the members—though the procedure is made up of a number of commitments by members to proceed in a certain way together.

States distribute tasks to members on the basis of their goals, which are a complex function of members' goals. To be sure, this function often turns members into adversaries and some members (whether for a good reasons or not) have more sway over the procedure than others. Decisions are not reached by anything remotely resembling consensus. Moreover, not all of the procedures are explicated and they often change in ad hoc ways. Sometimes, the state reaches internal impasses, indecisions, or different parts (members) of the state's agency pull the state in different directions. This, however, should not lead us to think that the state is not a unified agent: after all, individual agents regularly suffer from conflicting desires, *akrasia*, and so on, and we do not take this to mean that they are not unified agents. (Pettit 2003 makes similar observations about the parallels between individuals' internal conflicts and collectives' internal conflicts.)

The result of these processes within the state is a set of goals, a set of individual tasks for achieving those goals, and a distribution of the tasks among individuals, that is the decision not of one member—or of the conjunction of each member's independent processing—but of the members together, each acting within his or her part in the group's decision-making procedure. Thus, the will of the state is a complex function of the will of

<sup>94</sup> Under my terminology, then, Max Weber's famous definition of the state ("A state is that human community which (successfully) lays claim to the monopoly of legitimate physical violence within a certain territory" (Weber 1919[1946], ¶4)) gives a necessary, but not sufficient, condition. It should be noted that (d) does not require that all or even most other states recognise an entity as a state in order for that entity to count as a state. Similarly, Article 3 of the Montevideo Convention explicitly denies that recognition by other states is necessary. As Thomas D. Grant (1999, ch. 1) discusses, there are several competing theories of when an entity should be recognised as a sovereign state. Not all of these require that it is recognised as such by other states.

<sup>95</sup> Toni Erskine (2001) and Alexander Wendt (2004) agree, each using slightly different criteria for collective agency.

its members. The actions of those with roles under the state's decision-making procedure contribute to actions of the state when the role-holders act because of and within their role. These considerations combine to make states agents, under Chapter Three's account of collective agency.

Even if one denies the account of collective agency I gave in Chapter Three, there is a methodological reason for accepting that states are agents, and moral agents in particular. This is that we treat them as if they are so. It is common in moral and political philosophy to assume that there is a fact-of-the-matter about what characteristics an entity must have if it can bear moral responsibility, and that we can assign moral responsibility only to entities that have those characteristics. But if we look at the debates about what constitutes agency—such as metaphysical debates on free will, personal identity, and so on—we see that these debates take moral intuitions and moral practice as data (paradigmatically, Strawson 1962). Plausibly, we hold people responsible, and have reactive attitudes to them, just when we think they should (and therefore, think that they can) respond appropriately to criticisms of them. So rather than figuring out who the agents are before we assign responsibility to them, we might look at who we assign responsibility to and take that as evidence for those entities being agents.

Of course, we should not assume that all of our everyday responsibility-assigning practice is correct: as we saw in Chapter Three, we sometimes feel the urge to assign duties to random collections of people that clearly lack agency. Reflective equilibrium between moral intuitions and philosophical accounts of agency is almost certainly necessary. But the fact that metaphysical accounts of agency take moral practice seriously should perhaps lead moral and political philosophers to take it seriously, too. In this connection, consider reactions to the 2010 BP oil spill;<sup>96</sup> the Space Shuttle Challenger disaster;<sup>97</sup> or so on. If we are sure enough of our reactive attitudes in these cases, this gives us *prima facie* reason to impute the kind of agency necessary for moral responsibility to at least some collectives. Plausibly, states will be among them.

Let us continue, then, as though we have established that states are collectives, i.e., moral agents. How are we to draw lines around them? Who, in other words, constitutes their agency?

---

<sup>96</sup> In its final report on the incident, the Joint Investigation Team of the U.S. Coast Guard and the Bureau of Ocean Energy Management, Regulation and Enforcement made several attributions of responsibility to collectives. It said that BP "was ultimately responsible for conducting operations at Macondo in a way that endured the safety and protection of personnel, equipment, natural resources, and the environment"; that "Halliburton ... was responsibility for conducting the cement job, and, through its subsidiary (Sperry Sun), had certain responsibilities for monitoring the well" (BOEMRE 2011).

<sup>97</sup> E.g., at <http://history.nasa.gov/rogersrep/genindex.htm>

### ***6.2.2 Collective Members of Collectives***

There are different kinds of decision-making procedures that states might have: autocratic, oligarchic, monarchical, democratic, and so on. Whichever kind of decision-making procedure a state has, though, it will have a government. A state's government can be understood as the conjunction of the legislative (law-creating), executive (law-enforcing), and judicial (dispute resolving) branches of a state—though in a given case, a government might lack one of these or might have other functions as well. The government has the authority under the decision-making procedure to act on behalf of the state as a whole, to execute the state's decision-making procedure, and to interact, in their roles in the state's decision-making procedure, with the governments of other states.

The government is usually<sup>98</sup> a member of the state and, if so, it is a sub-collective of the state: that is, it is a member that is itself a collective. This means it has its own decision-making procedure (i.e., its own agency) *and* a defined role in the state's decision-making procedure. Collectives with collective members should not surprise us. For example, the Parliament of Australia is a collective, which is made up of two sub-collectives and an individual role bearer: the Senate, the House of Representatives, and the Queen. These three entities are aptly described as members of the collective that is the Parliament of Australia, even though two of them are themselves collectives and one of them is not. In the case of the Parliament of Australia, one of the members (the House of Representatives) is a collective that is itself made up of collectives. Specifically, the members of the governing administration and the members of the opposition respectively make up two mutually exclusive collectives. We can talk about both of these collectives as members of the House of Representatives, particularly when these two members respectively take unified “party lines” on particular issues. These two collectives are members of the House of Representatives, which is itself a member of the Parliament. Of course, like all collectives, at least some of the members of states are natural individuals and membership will always bottom out in natural individuals.

### ***6.2.3 Individual Members: False Negatives and False Positives?***

As with all collectives, we should understand each state as having as its members all those who figure in the conditions for collective agency. Assuming the account of collective agency I gave in Chapter Three is correct, then members are those who (a) are given roles

---

<sup>98</sup> I say “usually,” but in regard to the actual world I could say “always.” I say “usually,” though, because logically, we can imagine the government’s functions being contracted out.

by state's decision-making procedure; (b) are permitted by their role in the procedure to have influence on the procedure's results; (c) have an expectation that the procedure will produce aims, roles for achieving aims, and a division of aims; (d) have committed (perhaps tacitly or implicitly; perhaps only to themselves) to abide by the procedure's decisions; and (e) are able to receive sufficient information about the procedure's decisions for them to abide by it. Which individuals, then, are the members?

This is problematic: (a)–(e) seem, at first glance, to generate both false negatives and false positives. First, false negatives: an anarchist who is a citizen of a liberal democracy probably will not count as a member of that state, on my view. Plausibly, she will meet (a), (b), (c), and (e), but almost certainly will not (d). She is given the role of obeying the law, voting, paying taxes, and so on; she is permitted to vote, lobby parliament, petition, and so on; she expects for the state's procedures to produce aims, roles, and divisions of roles; and she can at any time ask the local Citizens' Advice Bureau all she needs to know to abide by the procedure. But plausibly, she has not committed to abide by the procedure's results. And laws that she obeys are not obeyed *because* they are produced by the procedure of a collective to which she has committed. This negative result seems incorrect.

Yet it gets something important right. There is a significant sense in which the anarchist is not part of (does not partly constitute) the state's agency. Her agency is not implicated when the state acts. This is consistent with her having a duty to be a member of the state, which she is failing to discharge. It is also consistent with her having a duty to obey the state's laws, which she might or might not be discharging. What this shows, I think, is something about which notion of "state membership" we are interested in, when we are in the specific business of determining who is implicated in the state's agency (and, more specifically, its moral agency).

Under the relevant notion of state membership, the people who count as members are perhaps more appropriately called the "decision-making members" than just "members," because we are interested in those members who have a share in the state's (moral) decision-making, and whose acts partly constitute acts of the state itself. These are the members I am primarily interested in, since I am interested in the state's duties. And the anarchist is not one of them, because she has not committed to abide by the procedure's results. (Again, this might in itself be a grave failing of hers.) So this seemingly false negative is, I suggest, not a false negative at all.

Yet my account also seems to generate false positives: people seem to count as "decision-making members" on my view, when they surely are not *any* kind of member. The population of, for example, Australia at any one time includes many permanent, short-term, and long-term residents. These include tourists, illegal immigrants, diplomats,

international students, permanent migrants without citizenship, and so on. It might seem that, according to (a)–(e), at least some of these individuals should be regarded as decision-making members. Take tourists. They are given the role by the decision-making procedure of “abiding by the legal rules the state produces.” This general role contains many sub-roles: abiding by road rules, passport and immigration rules, consumption tax rules, and so on. Tourists exercise *some* control over the procedure, in virtue of a conditional: if they withdraw their commitment (for example, by leaving), then the state can no longer give them the roles that it does. Assuming the state is not a failed state (in which case there would no collective procedure at all), tourists generally believe that the state is producing roughly the kinds of decisions, roles, and role distributions that it does. The tourists have committed—by virtue of signing arrivals documents—to abide by the state’s procedures. And they might (though not always will) be given enough information to perform their role satisfactorily.

While this may seem like a positive result on my account, the details will matter here. It is highly unlikely that individuals such as tourists, illegal immigrants, and so on have the *right kind* of influence on the procedure’s results. They are not permitted *by their role* to have influence on the procedure’s results. They can influence the procedure from the outside, like any (agential or non-agential) factor in an (individual or collective) agent’s environment can. They do this by leaving or arriving. But they do not exercise that control from the inside. This influence is not mandated by their role. Certainly, the government might *choose* to take the beliefs, desires, interests, and so on of these agents into account. It is within the scope of government officials’ roles that they have the discretion to choose to do so. But in that case, these agents’ influence on the procedure’s results depends upon the will of the other role-bearers (e.g., the government officials). The agents’ influence is not part of a role within the procedure.

These two examples—the apparent false negative and the apparent false positive—point to a more general conclusion. States with different procedures will have different types of decision-making members. We cannot assume, from the fact that a person lives in a territory (or even from the fact she holds the relevant passport) that she is a decision-making member there. This matters, as it will have implications for how we distribute to individuals the duties of real-life states. Two imagined cases will help to concretise who counts as a decision-making member of a state, that is, a member for the purpose of that state having and discharging a duty.

#### 6.2.4 Autochria

Imagine a territory—call it Autochria—in which authority over laws, policies, and interactions with other states is held primarily by one individual, though this person has a small number of official advisors. The decision-making procedure of Autochria is simple: it is discussion between the leader and her advisors. The advisors have the official role within the procedure of voicing their opinions to the leader. If the leader does not listen to the advisors, the procedure allows that the advisors vote out the leader. This ensures that the advisors' opinions do influence the decision-making procedure's outputs, in virtue of their role as advisors. But no other members of the population have a role in the mechanism that gives them any influence over the mechanism's outputs.

In Autochria, the only decision-making members of the state are the leader and her advisors. That is, the state is identical to the government. The advisors' level of influence over the decisions is not precisely determined by the procedure. But it is determined that they will have some. By contrast, many of the people we would naturally call “citizens” of Autochria are not decision-making members of the state on my analysis, as they do not have any influence over the mechanisms of the procedure.<sup>99</sup> The majority of citizens meet conditions (a), (c), (d), and (e) for decision-making membership. Their role is to obey the laws that the state produces. Suppose that they have tacitly committed to act within this role. They have an expectation, commonly known to be held by the relevant others, that enough other role-bearers will abide by the procedure. They are able to receive sufficient information about the procedure's outputs for them to abide by it. Yet they lack a role in the mechanism that gives them influence over the mechanism's outputs. Rather, Autochria must treat citizens as features of the environment when it is discharging its duties. These are features of the environment that are manipulatable and intentionally responsive, but they are nonetheless not part of the collective agent itself.

Of course, this does not give Autochria moral free reign to disregard its citizens' interests. Autochria will almost certainly have duties *to* these non-member citizens. But these duties will not be constituted by citizens' duties: in no cases will citizens be the ones who must use the Autochria's decision-making procedure in order to discharge these duties. We are here interested in the question of which individuals might incur duties to use the state's procedure to discharge the state's duties.

---

<sup>99</sup> We might think that surely the citizens could rise up against the autocrat, and that this possibility must influence her decisions. In this way, one might say that the citizens indirectly and externally influence the outputs of the decision-making procedure. But if there is no legal right of protest, petition, etc., then this is not done in virtue of their role.

### 6.2.5 Democracia

Imagine a second state, Democracia. In Democracia's decision-making procedure, the vast majority of the permanent members of the population possess a particular set of rights and obligations. The obligations of these people are to pay taxes and obey the laws the government makes, and the rights are to vote in periodic elections. The decision-making procedure is such that the government is chosen through a complex, but broadly majoritarian, function of these people's votes. Not all permanent members of the population have these rights and obligations, however. These include permanent migrants who have not gone through certain naturalization processes, persons under the age of eighteen, and persons in prison. But the people that do have these rights and obligations are decision-making members of Democracia. That is, (a)–(e) apply to them. Their rights and duties attach to their role; their votes constitute their role-based input into the procedure; they have (almost all) committed to abide by the outputs of the procedure; they expect others to abide; and they receive enough information about their role in order to perform it.

This is consistent with these people sometimes judging a particular law or policy to be unjust or otherwise undesirable. The important point is that they commit to abide by the procedure, not that they agree with the substantive results of that procedure. Even civil disobedients will still count as members if they have committed to abide by the outputs in general, despite not abiding by particular ones.

One might object that, in Democracia, the people who merely have voting rights and law-obeying obligations are not properly construed as members. If the state is large enough, then one member's vote will have a very limited effect on the outputs of the collective's decision-making procedure. So perhaps we should rather identify the members of the state with the members of the government—those who meet (a)–(e) and have a *tangible* effect on the collective decision-making procedure's outputs, or some such. However, there is no principled way to draw the line of “sufficiently consequential input to count as a member” in between ordinary citizens and, for example, clerks (or human resource managers, or...) at government departments. Both of these roles are similarly inconsequential for the outputs of the state's decision-making procedure.

One might reply that, while “being a member of the state” is a vague concept, ordinary citizens clearly do not fall under it, prime ministers clearly do, and some government workers (clerks, human resource managers...) fall in the vague area. I deny this intuition. All collectives have more and less “important” members. A given member's role in a state might be as insignificant as “voting in the general election,” or as significant as “chairing the Cabinet.” Moreover, at least some governments do, in fact, represent their

ordinary citizens—or at least try to or claim to.<sup>100</sup> It seems difficult to account for this representational function of (some) governments unless both the governments and the ordinary citizens are subject to an over-arching decision-making procedure.

If I'm right about who the decision-making members are in Autochria and Democracia, this has important implications for how we should conceive of the duties of real-world states, and for how we identify the individuals to which those duties can be distributed. When a given state has a duty, which agents have distributed duties will depend entirely on that state's decision-making procedure. Only agents who have a role in the state that affords them positive influence under the state's decision-making procedure can bear distributed duties, since only these agents partly constitute the state's moral agency. This helps us to understand the sense in which ordinary citizens of democratic states are implicated in their states' agency (and bear distributed duties when it has duties), while ordinary citizens of non-democratic states are not. (Similar conclusions are reached by Pasternak (forthcoming) and, using different arguments from those given here, by Collins and Lawford-Smith (2013).)

### **6.2.6 The Need for Non-state Collectives**

States are not the be-all and end-all of large-scale real-world collective dependence-based duties. There are three kinds of reason for this. The first kind of reason relates to the capacities of states. Even if every state complies with its duties, there still might be some important interests that states are insufficiently capable of fulfilling, or that states are not best-placed to fulfil. This might be because of states' budgets, international commitments, susceptibility to corruption, and so on. Or there might be some duties that, by their nature, can be discharged only through a collective where members' performance of their role-based duties is voluntary. Perhaps this is because the duty should be discharged from a certain altruistic motivation—as is the case with duties to have caring attitudes. In these cases, if individuals are coerced into being members of the collective or are coerced to take certain roles in the collective (as the state might coerce them), then their resultant actions or attitudes won't count as discharging that duty. This is not to deny that role-bearers within the state cannot, or even often do not, have altruistic motives. It is just to suggest that *some* of the time, in *some* cases, the state cannot conjure up such motivations. In these

---

<sup>100</sup> Along these lines, David Runciman (2012, 67–68) claims that “[r]epresentation is a key element, perhaps even the key element, in the legitimacy of the modern state...”

cases, non-governmental organisations might be the collective best-placed to ensure that members have this kind of motivation, given that membership is voluntary.<sup>101</sup>

Second, states may be unwilling—even if they are best-placed. States are generally reluctant to do what is not in their interests, especially when it comes to actions beyond their borders. In particular, democratic states are swayed into discharging whatever duties citizens conceive of themselves as having—and citizens are not always correct about what these duties are. This will generate a dependency duty for the next-best-placed collective, given that, when we are assessing whether one agent has a dependency duty, we are realistic about the prospects of others abiding by their duties.

Third, and crucially for the remainder of this thesis, sometimes states can realise more value if they coordinate with each other in the fulfilment of important interests. That is, states may fail to meet the Dependency Principle for taking particular measures because the value of those measures is rendered negative by their opportunity costs—perhaps the opportunity costs for the states themselves, but more compellingly the opportunity costs for dependents. Precisely, the opportunity cost at issue is the loss of the opportunity for the state to *take responsive measures*—in combination with other agents, usually themselves states—to fulfil the relevant interests.

Consider, for example, asylum seekers' interests in obtaining asylum. Each state is capable of doing something (granting asylum) for some of these people. But the most valuable way (regarding dependents and agents) for these interests to be fulfilled is not for one state to unilaterally take measures that help just a few asylum seekers. Instead, the most efficacious way for those individuals' interests to be fulfilled is for the interest to be *packaged*. According to the Dependency Principle, the “most efficacious” measure for fulfilling some interest is the measure that *either* “has the highest likelihood of fulfilling that interest,” *or* “of the measures that are tied for highest likelihood of fulfilling that interest, best fulfils *other* important interests.” While any one state could grant asylum to any one asylum seeker, often the most valuable way for a state to grant that asylum will be for it to do so as part of a *global system* of asylum. That global system is the most efficacious way to fulfil the interests of the one asylum seeker, because it has the knock-on effect of fulfilling other asylum seekers' important interest in other states having committed granting asylum.

But in order to explain any one state's duty to take measures to set up a global asylum system, we cannot appeal to *that* state's being best-placed to set up the system. This is because no state is best-placed—on its own—to set up such a system. (This is analogous

---

<sup>101</sup> Goodin (2003) discusses the “motivational distinctness” of nonprofit NGOs. On the lack of entailment (in either direction) between an organization's type (for-profit, state, or nonprofit) and the organization's potential for voluntary altruism, see Alexander 1987.

to Chapter Three's discussion of the drowning swimmer.) We must instead appeal to each state's duty to take responsive steps to set up such a system. We must appeal to their *coordination duties* to do so. As it happens, international politics is (slightly) ahead of the game here: there already exist collectives—of which states are members—that can bear dependency duties to fulfil the interests with respect to which states would (if not for these collectives' existence) have coordination duties. These collectives are inter-governmental organisations. If inter-governmental organisations really are collectives, then they can bear dependency duties that block the existence of some of states' coordination duties. That is, inter-governmental organisations might be able to render (some of) states' would-be coordination duties redundant. If this is right, then instead of coordination duties, states will have duties to *act within their roles* in inter-governmental organisations with a view to the fulfilment of those organisations' duties. States will have duties *qua* members of a collective, not duties *qua* ad hoc (but potentially responsive) agents. As a precursor to discussing R2P, it is worth considering whether this is the right analysis.

## 6.3 Inter-governmental Organizations

### 6.3.1 The General Picture

Do inter-governmental organisations (IGOs) have the agency necessary to bear dependency duties, thus blocking the need for (some) coordination duties of states? In order to answer this, we need to be clearer about what IGOs are. Very roughly, an IGO is a collective, created by agreement between states, of which the members are states or organs of states.<sup>102</sup> IGOs generally have a permanent secretariat, staff, and/or headquarters. There are numerous IGOs. They might aim to be maximally inclusive, such as the United Nations (UN) or International Criminal Police Organization. They might select member regionally (e.g. the African Union or European Union), or historically or politically (e.g. the Commonwealth of Nations or Organisation Internationale de la Francophonie).

If IGOs are really collectives (i.e., agents) then there is nothing problematic about the fact that the members of IGOs are themselves collectives, just as long as the members are agents—that is, as long as the members are proper collectives rather than random

---

<sup>102</sup> The International Law Commission's 2011 Draft Articles on the Responsibility of International Organizations gives a similar definition of "international organization," saying that "'[i]nternational organization' means an organization established by a treaty or other instrument governed by international law and possessing its own legal personality. International organizations may include as members, in addition to States, other entities" (ILC 2011, 2). This definition is somewhat broader than the gloss I gave (allowing as it does for non-state members and non-treaty formations), but we should not assume this to allow that, for example, international corporations or international non-governmental organisations are included in this category. IGOs are essentially state-based entities.

aggregates or merely coordinated groups. Any collective with its own decision-making procedure is an *agent* that can make decisions, have goals, and perform tasks distinctly from its members. As long as the members of IGOs are agents in this sense, the question of whether the IGOs themselves are collectives can be answered by using the same criteria we used in Chapter Three—and for states, above—for determining whether a group of agents is a collective.

IGOs meet these criteria. They have members, which are agents (in this case, states). These members are given roles by the procedure, usually in the form of a number of rights and duties listed in the IGO's founding and subsequent treaties. These membership-based rights typically include a right for the role-bearer to have influence over the collective's decisions, under the procedure.<sup>103</sup> The fact that the procedure takes inputs from a number of states entails that the IGO's decision-making procedure is distinct from those of its members. The procedure generates actions the IGO will perform, the individual roles that are jointly sufficient for those actions, and an allocation of these roles. At the most general level, the actions include things like the maintenance of international peace and security,<sup>104</sup> the facilitation of cooperation amongst states' law enforcers,<sup>105</sup> and so on; at a more specific level, the actions might be calling for restraint in civil wars,<sup>106</sup> or training states' criminal investigators.<sup>107</sup> In joining the IGO, each member explicitly agrees to abide by the procedure's decisions (or at least, to abide by most of them: states can enter "reservations" when signing treaties, but when the reservations arise they can be treated in the way civil disobedients were in the above discussion of states). And the procedure's results are communicated to members in a way that is sufficient for them to perform their role.

The attribution of agency to these organisations is backed up by law. In his widely-used textbook on public international law, James Crawford (2012, 169–70) states that possession of "legal personality"—personhood or agency under the law—is a necessary condition for an entity to be an international organisation. He notes that numerous

---

<sup>103</sup> For example, in the United Nations Charter (ch. IV, art. 10): "The General Assembly may discuss any questions or any matters within the scope of the present Charter or relating to the powers and functions of any organs provided for in the present Charter, and, ... may make recommendations to the Members of the United Nations or to the Security Council or to both on any such questions or matters." <<http://www.un.org/en/documents/charter/chapter4.shtml>>

<sup>104</sup> The most general aim of the United Nations, <<http://www.un.org/en/documents/charter/chapter1.shtml>>.

<sup>105</sup> The most general aim of the International Criminal Police Organization, <<http://www.interpol.int/About-INTERPOL/Overview>>.

<sup>106</sup> Again, the United Nations, <<http://www.un.org/apps/news/story.asp?NewsID=45126&Cr=Libya&Cr1=#.UbWg-YLT5jI>>.

<sup>107</sup> Again, the International Criminal Police Organization, <<http://www.interpol.int/en/News-and-media/News-media-releases/2013/N20130605bis>>.

entities—most prominently, the United Nations and its organs—are granted such recognition. According to Crawford, legal debates about whether (for example) the North Atlantic Treaty Organization or European Union is an international organisation are precisely debates about whether the organisation has a “separate personality” from that of its members (2012, 169–70). It thus seems that, according to international lawyers, IGOs are agents *by definition*. Of course, if lawyers have bad criteria for deciding whether or not an entity has legal personality, then the fact that an entity is legally recognised as an IGO does not philosophically prove that it is an agent. Yet it is doubtful that the numerous and varied instruments for granting an agent legal personality (as described in Crawford 2012, 172–4) could go so very wrong in formulating or applying their recognition mechanisms. The fact that IGOs are recognised as agents in law at least gives us some reason to see them as such.

Given that there are already IGOs, these group agents can directly incur dependency duties. Indeed, many IGOs seem already to be discharging something roughly like dependency duties, and seem to have been set up for this purpose. We can say that their founders were fulfilling coordination duties in founding these IGOs—even though they might not have viewed their actions in these terms. Consider, for example, the IGOs that are charged with enforcing various declarations, conventions, and covenants.<sup>108</sup> These IGOs’ founders coordinated in order to bring about the fulfilment of important interests, by setting up a collective that does not fulfil the interests directly, but rather formulates, advocates, and commits members to various declarations, conventions, and covenants that instruct members to do so. Once this step was achieved, the IGO was able to organise itself around the goal of seeing to it that members stand by the individual commitments that they made to abide by the group’s decision-making procedure.

Often, the IGO is needed to bring about the fulfilment of important interests because each member requires assurance that other members will also perform some task before that member’s performance of its own measures towards this end will be efficacious, and therefore generate positive expected value regarding the agents and dependents involved.

Importantly, the criterion that agents must have a decision-making procedure forbids us from recognising as “IGOs” such aggregates as “the international community”

---

<sup>108</sup> The declarations, conventions, and covenants include, for example, the Universal Declaration of Human Rights (1948); the Convention on the Prevention and Punishment of Genocide (1948); the International Covenant on Civil and Political Rights and the International Covenant on Economic, Social and Cultural Rights (both adopted in 1966); the Elimination of all Forms of Racial Discrimination (1969); the Convention on the Elimination of all Forms of Discrimination against Women (1981); and the Convention on the Rights of the Child (1990). The IGOs that are charged with enforcing these standards include the Human Rights Committee; the United Nations High Commission for Human Rights; the Human Rights Council; the International Criminal Court; and various ad hoc tribunals. See Nickel 2010 for a description of these bodies.

or “the society of states”—even if, in some bizarre future, the law were to recognise them as such.<sup>109</sup> These aggregates are not agents, so cannot be duty-bearers. (Indeed, it’s questionable who even belongs to them: are international corporations in the international community? Does North Korea belong to the society of states?) It is worth pointing this out, as international political agents sometimes use language in a way that suggests otherwise.<sup>110</sup>

### ***6.3.2 IGOs as Real Collectives***

One might take issue with my characterisation of IGOs as collectives. Yes, state members are bound by the decision-making procedures of IGOs, so are obliged to follow the role instructions that they are given by the IGO’s decision-making procedure (unless the state has entered a “reservation” on that point, upon acceding to the treaty establishing the IGO). This includes abiding by the conventions that the IGO endorses on behalf of its members. But surely, one might think, there is a loophole here: given that IGOs are run by their members (states), IGOs can only distribute to states those roles that states have given the IGO decision-making procedure the power to distribute—and IGOs can only distribute those roles in the way that states permit through the formulation of the decision-making procedure. IGOs are, perhaps paradoxically, both controlled by, and controlling of, their members (Abbott et al 2010, esp. 2). Yes, IGOs can and sometimes do provide “warrants,” “approval” and “legitimacy” to states’ actions (as Harbour (2004, 67) argues in defence of their status as moral agents). But they are not the ones executing the actions and their instructions are not always taken as authoritative. As even Crawford notes, “[u]nlike states, international organizations do not possess general competence: they may only exercise those powers expressly or impliedly bestowed upon them” (2012, 184). We should therefore be sceptical of the claim that they have agency of their own. Or so the thought goes.

---

<sup>109</sup> Although we might think that some IGOs somehow *loosely* represent, or *tend to* reflect the views of, the set of agents in these ad-hoc groups. Chris Brown, for example, notes that states often act as though the United Nations Security Council has the capacity to act on behalf of international society: “[t]he attribution of a key role to the Security Council as the *bearer of agency on behalf of international society* [at least in respect of things covered by the UN Charter] rests on the UN Charter, the behavior of its members, including the most powerful of them, and, to an extent, on public perception...” (Brown 2001, 91, emphasis added). This does not make “international society” an agent.

<sup>110</sup> To give just one representative and recent example, in October 2012 United Nations Secretary General Ban Ki-moon told an audience at Drake University that “[t]he international community failed to protect thousands of Bosnian Muslim men and boys from slaughter. ... The international community has a moral responsibility, a political duty and a humanitarian obligation to stop the bloodbath and find peace for the people of Syria.” (UNDPI 2012b). We will see further examples of this kind of rhetoric in Chapter Seven.

However, to some extent or other, this is a problem for all collectives. Collectives are controlled by their members or by some subset of their members. And all collectives can be disbanded or have their procedures overridden if enough members (or the right ones) decide to take action to make it so. Any member of any collective might withdraw from the collective—with both the collective and the member having to face the consequences. The collective's agency is constituted by its distinct decision-making procedure. This is all that is needed to imbue it with agency.

Yet the objector might have got something right here: IGOs' *capacities* might be very little over and above the sum of their members' capacities when those members are acting in a mutually responsive (but IGO-independent) way. Thus while we should, I suggest, maintain that IGOs have *agency*, the objector's point might have important implications for IGOs' *duties*, in particular their dependence-based duties (based as they are on agents' capacities). Perhaps IGOs do not have many dependency duties, because all important interests could be fulfilled by states acting independently of IGOs. If states are, through mutual responsiveness, better-placed to fulfil all the interests that IGOs might fulfil, then there is no point in discussing IGOs in a thesis on dependence-based duties.

Indeed, scholars of international law and international relations tend to agree. Kenneth W. Abbott et al point out that IGOs

lack direct access to private actors or other targets of regulation (because they do not have either the authority or regulatory capacity), lack adequate decision-making capacity (because they operate under consensus-oriented decision rules in a context of heterogeneous preferences), and/or lack adequate monitoring and enforcement capacity (because they have limited authority, revenue and administrative staff). (Abbott et al 2010, 1)

Abbott and Duncan Snidal (2009) suggest that IGOs have independent capacities only insofar as the IGO can bring in third parties (such as NGOs or private actors) to act as intermediaries between IGOs and those actors (usually states) that IGOs wish to influence. The intermediaries are likely to have information, access to actors, monitoring capacities, or sanctioning capacities that the IGO lacks. In short, IGOs lack the capacity for efficacious action or influence without the voluntary assistance of other, more autonomous actors. Similarly, Frances Harbour (2004) looked at six cases of humanitarian intervention and found that in none did states base their decisions on IGOs' evaluations. And Gareth Evans points out that “[s]tates remain, for better or worse, and will be for the foreseeable future, the primary actors in international affairs. Intergovernmental organizations can only decide

if their member states agree and can only act if their member states deliver, and nongovernmental organizations can only be influential in their advice or effective in their program delivery if individual states listen and allow them to act" (2008, 196).

Yet these statements, while perfectly true, neglect some important facts about IGOs' capacities. First, IGOs often *make possible* the very multilateral responsiveness that is necessary for states to reliably bring about desired (or desirable) states of affairs. Often, they do this by acting as *fora* within which states can coordinate their actions. This is rather a weak role, with states using their status as IGO members to reach decisions that are not themselves decisions *of the IGO*. For example, states might use IGOs' physical and administrative apparatuses to provide each other with the requisite assurance that they will  $\varphi$  if others  $\varphi$ . This is not quite the same as states' agreeing on how they will make future decisions (so it is not using the IGO to form a collective agent), but it does allow states to produce outcomes—such as “multilateral  $\varphi$ -ing”—that they could not produce without IGOs. The fact that IGOs are used in this way points towards a special capacity *of IGOs*: the capacity to *encourage* and *facilitate* states' use of IGOs' apparatuses in this way.

Second, such inter-state coordination will sometimes require that IGOs are not just the *forum* for decision-making, but that they are the decision-*maker* itself. Some outcomes require that multiple states agree on how they will decide before the time for decision comes. This is more than just an agreement or act of mutual responsiveness: it is the setting up of a group decision-making procedure to which they are all committed, which will come into effect under conditions on which they all agree. It is the setting up, in fact, of an IGO, which will be the decision-maker when the time comes to decide.

A third source of IGOs' special capacities derives from the fact that any one state cannot be guaranteed that the IGO will give them the tasks that the state would prefer. The quotations above might imply that *each* state is more powerful than the IGO. In fact, states *together* determine the scope of IGOs' powers. The power of any one state—even a powerful state like the United States—is therefore much weaker than that of the IGO, within the IGO's scope of decision-making and assuming that the state abides by the IGO's decisions. IGOs' decision-making procedures are a result of complex and lengthy negotiations between members, with the result that members will often receive roles that require them to perform actions that they would not perform in the absence of IGOs. Of course, more powerful states are much more likely to be designated a role that they want—but they are still not guaranteed to have their preferred action permitted or their least preferred action not demanded (Woods 2003; Grigorescu 2007, 296–7; Steffek 2010). The capacity to distribute roles against the will of states gives IGOs the capacity to produce some outcomes that would be very unlikely to result from states' mutual responsiveness.

Of course, it remains true that many IGOs are controlled by a few powerful members, which use the IGO to further their own interests rather than using it to discharge states' collective duties which IGOs incur through their creation by states. One example is the International Monetary Fund (IMF), where each state has voting rights that are "based broadly on its relative position in the world economy" (IMF 2013). Another is the United Nations General Assembly. There, decisions and motions are approved usually by "consensus" within each regional group. Members rarely oppose the majority opinion within their group, because the benefits that a state obtains within the UN system are determined by their group (Wedgwood 2009). Yet another is the United Nations Security Council (UNSC), whose decision-making procedure confers veto power upon the five permanent members. (I will discuss the UNSC further in Chapter Eight.) But these examples do not refute the more general point that IGOs have agency and capacities—and, therefore, quite possibly *duties*—to realise outcomes that their members could not realise in the absence of the IGO.

## 6.4 Conclusion

My comments in this chapter have been inevitably brief. Nonetheless, I hope that they have gone some way towards assuaging any concerns about the applicability of Part I's theory to the unwieldy large-scale collectives that are states and IGOs. This step is crucial if I am to use my account of dependence-based duties to precisify, unify, and explain an important doctrine within international ethics: the Responsibility to Protect.

# Chapter Seven:

## **Responsibility to Protect: Claims and Potential Explanations**

### **7.1 Introduction**

In his report *We the Peoples*, former United Nations Secretary-General (UNSG) Kofi Annan posed the question: “If humanitarian intervention is indeed an unacceptable assault on sovereignty, how should we respond to a Rwanda, to a Srebrenica—to gross and systematic violations of human rights that offend every precept of our common humanity?” (Annan 2000, 48) In asking this, Annan drew attention to a deep tension in the international system. On the one hand, it has long been recognised that each state has a duty not to intervene in other states’ internal affairs. On the other hand, there is a growing conviction that states have duties to protect individuals from certain horrors—even, on occasion, when those individuals live in another state and protection would require intervention in that state’s internal affairs. The governing principles of international relations—not to mention the lives of countless human beings—depend on how we resolve the tension between these two propositions.

The emerging doctrine of Responsibility to Protect (R2P) aims to resolve this tension between non-intervention and human protection. In short, R2P prioritises the latter: it asserts that the duty to protect individuals from certain egregious harms trumps the duty not to intervene in states, where the two duties conflict. It does this by asserting that each state—in this context, synonymous with “government”—has a duty to protect its own population from certain egregious harms. States and other international agents (especially IGOs) have duties to support each state in discharging that protection duty. But if a state is either unable or unwilling to protect its population, then other international agents acquire duties to protect that population. As a last resort, fulfilling this duty may require military intervention.

R2P is a doctrine of increasing international prominence. Among international actors—for example, delegates who congregate at the United Nations General Assembly (UNGA)—there is broad support for it, at least in its non-military guise (ICRtoP 2009; GCR2P 2009). It has been increasingly evoked by the United Nations Security Council (UNSC) in recent years, with the UNSC using R2P language in resolutions addressing

events in Libya, Yemen, and Syria, among others (UNDPI 2012a).<sup>111</sup> In September 2012, United Nations Secretary General (UNSG) Ban Ki-Moon told the UNGA that, over the previous year-and-a-half, R2P had been “front and centre as never before” in the affairs of the UN (UNDPI 2012a).

Moreover, something approximating R2P seems to have public support. According to a global opinion poll run by The Chicago Council on Global Affairs and WorldPublicOpinion.org (2007, 3–4), an average of 73% of respondents thought that the UNSC should have the “right to authorize the use of military force to prevent severe human rights violations such as genocide,” and an average of 57% said the UNSC had a “responsibility” to authorise such force.<sup>112</sup>

The R2P is, then, clearly of great political importance. It might be used—and has been used—to justify extremely large-scale and significant actions of states. However, the R2P is a very diffuse collection of related normative propositions. The doctrine is diffuse in that it cannot be precisely identified with any one published variation on those propositions—especially since these variations conflict, and are each under-specified and under-theorised. In particular, the *moral basis* of R2P is largely unexamined by its canonical formulations. When it is examined, the results are vague and unconvincing. Even academic commentators have largely shied away from reflecting on R2P’s basis, preferring to examine specific applications of the doctrine—particularly military interventions.<sup>113</sup> When R2P’s basis is considered, both the canon and the commentators appeal to clichés about the obligations of the “international community” or to unspecified platitudes about fulfilling human rights.

One might think that we can get by without examining R2P’s basis, when it comes to actually implementing the doctrine. Perhaps for practical purposes we can interpret the doctrine thus: states have duties to protect their populations, and when states fail, every agent acquires duty to do what he or she can to prevent harms to that population at not

---

<sup>111</sup> Though for an argument that these resolutions, in particular those regarding Libya, do not indicate that UNSC members in fact support R2P, see Hehir 2012, 12–20.

<sup>112</sup> These numbers result from averaging the responses from 10–12 countries, with a variety of economic, social, and political conditions, which together hold 56% of the world’s population. It should be noted that the responsibility to *authorise* force should not necessarily be read as synonymous with the responsibility to actually *use* force.

<sup>113</sup> There are a few exceptions to this, such as Welsh and Banda (2010) and Tan (2005). These authors’ treatments of the doctrine’s basis are brief, but I will mention them below. Of course, academics have long-theorised about the basis of humanitarian intervention obligations (e.g., McMahon 1996; Buchanan 1999; Bagnoli 2004; Pattison 2010). But they have not tied these obligations specifically to the R2P, and, as we shall see, R2P includes more than just humanitarian intervention.

too high a cost.<sup>114</sup> This might be fine as far as it goes. But the idea that the latter duties are held by *all* moral agents gives us little traction on the real problems. Regarding any particular harm to any particular population, it will often be best if only *one* agent (or small group of agents) does the protecting, when the home state fails. Such a division of labour allows protectors to focus on one population at a time, which is helpful if there is a decreasing marginal utility in having more protectors, or if information is lumpy, or if the investment of one protector makes that protector more efficient, or similar. Additionally, it will often be the case that not all potential protectors are up to the task. We want to know how we—or, more importantly, the UNSC—should go about authorising particular protectors.

We might say that all agents have a duty to protect all populations (insofar as they can) *until* someone else starts to do so effectively. But this would allow that the intuitively least appropriate (but still capable) agent might step in to protect, when other agents would be far preferable. This is morally undesirable—we have reason to want to pick out the most appropriate agent (or set of agents) as the bearers of these duties so that the moral impetus to protect is recognised as on *them*, and on others only if these “primary” duty-bearers fail. That is, we want to know *how* to distribute particular duties, to particular agents, for protection of particular populations.<sup>115</sup> This is because it is morally better to assign protection duties to the most morally appropriate agent, and to the next most morally appropriate if the first agent fails.

(This is not necessarily to say that the latter agents are *forbidden* from acting unless the former agent fails. It’s just that their doing so would be supererogatory, and the former agent would have done something wrong in not acting before they did. This is also not necessarily to deny that other agents have duties to *compensate* the agent who bears and discharges the protection duty, so that, in the end, the costs of protection are distributed somewhat equally among agents—the problem I am addressing is that of who should take the *initial action* of actually doing the protecting.)

There are also practical reasons for wanting to pick out particular duty bearers. Academic supporters of R2P are in very wide agreement that the duties to step in when a home state fails—in particular, where those duties require intervention—need to be

---

<sup>114</sup> This view is very closely reminiscent of Goodin’s (1988, 678, 684) view that strong obligations to compatriots are “merely devices whereby the moral community’s general duties get assigned to particular agents,” and that “where somebody is left without a protector he becomes the ‘residual responsibility’ of all.”

<sup>115</sup> This makes our problem quite a different one from that addressed by Goodin (1988), who concluded that states’ special responsibilities for their citizens “derive wholly from the fact that they *were* appointed [i.e., appointed as the bearer of a special responsibility that derives from the general responsibility all agents have to protect that population], and not at all from any facts about why they *were* appointed.” (Goodin 1988, 680) Our question is exactly what the mechanism of appointment—the relevant “facts”—should be.

specifically assigned to some agent or small group of agents in each instance, if they have a hope of *actually being fulfilled* in the world of realpolitik (Bellamy 2008, 429; G. Evans 2008, 206; Tan 2005, 88; Welsh 2007, 43–4; Welsh and Banda 2010, 219). This is because international political actors—primarily, states and coalitions of states—are quick to use others’ “inhumanitarian non-intervention” (as Simon Chesterman (2003) dubbed it) as an excuse for their own inhumanitarian non-intervention. If we could posit unified, precise, explanatory criteria for distributing duties of intervention, then the inhumanitarian non-intervention of some could not so easily serve as an excuse for the inhumanitarian non-intervention of all. This would be because not all agents would have the same type or strength of reason for action: one agent would be picked out by some criteria that does not apply to anyone else (or that applies to someone else only if the first agent defaults or fails). This agent could not so easily use others’ reticence to excuse their own.

One key problem with R2P is thus: when a state requires support in protecting its population, or when support would not be enough and intervention is required, how should we determine which international agent (or, potentially, group of agents) has the primary duty to act—which criteria should determine to whom we look first? As we shall see, the canonical formulations of R2P do not answer this question. To answer it, we need some criteria for picking out the duty-bearing agent. Clarifying the moral basis of R2P duties will take us some way to establishing these criteria. Clarifying the doctrine’s basis will serve other purposes, too. It will help to demarcate the limits of R2P duties, particularly R2P duties that require agents to act across state borders. An adequate explanation of R2P should not make international duties of human protection counter-intuitively common, counter-intuitively rare, too all-inclusive, or too costly relative to their benefits. Finally, the doctrine’s basis—if it is to be a unified basis for the whole of R2P—should also explain why the “stepping in” duties come second, that is, why states first have the duty to protect *their own* populations, with others acquiring duties only if the home state fails.

An adequate basis for R2P would thus help us answer three questions: first, why (morally speaking) the home state should have the primary duty; second, how to determine which particular agent(s) bear duties when the home state fails, and why this method of determination is the right one; and third, what the limits of these duties are. Over this chapter and the next, I develop my answers to these questions. The present chapter examines and rejects the answers found in the canonical formulations of R2P and in the surrounding academic literature. The next chapter gives my own positive proposal.

To that end, §7.2 examines the three canonical formulations of R2P. These three formulations are found in the 2001 report of the International Commission on Intervention and State Sovereignty (ICISS), the outcome document from the 2005 World

Summit of the UNGA, and a 2009 report of UNSG Ban Ki-Moon. By dissecting these documents, we will be able to fully appreciate the inadequacy of the canonical attempts to describe the basis of the doctrine and the international duties it implies, while also identifying the three core claims that constitute R2P.

In §7.3 I consider, and express some doubts about, possible ways of unifying, precisifying, and explaining these claims. These possibilities are sometimes hinted at in the canonical formulations, are sometimes suggested in the academic literature surrounding those formulations, and are sometimes derived from principles evoked by political philosophers and normative ethicists more generally. All of these suggestions, I will argue, are lacking as fully explanatory necessary and sufficient conditions for R2P duties. This chapter's negative project will thus establish what is problematic about the R2P canon and about some possible ways of improving on the canon. With this negative project completed, the way will be cleared for the positive project of Chapter Eight.

It might be worthwhile here to point out the senses in which this argument is and is not parallel to that of Part II. Compared to R2P, the academic literature on care ethics is huge. It thus took a full chapter to extract and disentangle the core claims of care ethics. In the case of R2P, however, there are relatively few canonical statements and the core claims are relatively straightforward (though, as we shall see, they underdetermine their own justification and precisification). The extraction and disentangling of the core claims will therefore be a briefer task for R2P than it was for care ethics. This task will be completed entirely in §7.2. The problem with R2P lies in the underdetermination of the *justification* of these claims. This is problematic since the justification feeds back in to the claims themselves, to affect the precise demarcation of the duties the claims entail. Thus, having established the (non-unified, non-precise, and non-justified) core claims in §7.2, §7.3 will concern itself with some potential unifying justifications of those claims. As we shall see, in rejecting these justifications we will simultaneously be rejecting some possible precisifications of the claims. To that extent, the present chapter ultimately does for R2P what Chapter Four did for care ethics: it philosophically critiques a particular set of normative claims, and rejects some ways of understanding those claims. Then, as Chapter Five did for care ethics, Chapter Eight will give my positive proposal for unifying, justifying, and precisifying the claims of R2P.

## 7.2 The Canonical Formulations of the Three Claims

### 7.2.1 ICISS 2001

In response to Annan's (1999a, 1999b, 2000) characterisation of the tension between state sovereignty and human protection (outlined above), the Canadian government commissioned the International Commission on Intervention and State Sovereignty (ICISS). This twelve-member panel of political and legal experts was tasked with researching and writing a report on the issues Annan had raised, through consultations with governments, NGOs, lobby groups, and other stakeholders. In 2001, it published its findings under the title *The Responsibility to Protect* (ICISS 2001).<sup>116</sup> Like others (G. Evans 2008; Bellamy 2009), I will take the ICISS's 2001 report as the primary statement of R2P.<sup>117</sup> According to the ICISS, the two “basic principles” of R2P are:

- A. State sovereignty implies responsibility, and the primary responsibility for the protection of its people lies with the state itself.
- B. Where a population is suffering serious harm, as a result of internal war, insurgency, repression or state failure, and the state in question is unwilling or unable to halt or avert it, the principle of non-intervention yields to the international responsibility to protect. (2001, XI)

Principle A outlines what I will call the “domestic duty”; principle B outlines what I will call the “intervention duty.” The ICISS divides both of these duties into three components: “preventing” the relevant harms; “reacting” if those harms occur; and “rebuilding” if necessary after the reactionary component. Of these three, the ICISS claims that prevention should be the top priority for duty-bearers (2001, XI), though the reactive component receives far more discussion in the report.<sup>118</sup>

Though it is not listed as a basic principle, the ICISS posits a third duty that I will call the “support duty”. This is the duty of international agents to support states in discharging the domestic duty. The support duty does not presuppose that the supported states are “unwilling and unable,” or that the principle of non-intervention “yield” (that is,

---

<sup>116</sup> This report should not be identified with the doctrine of R2P, which *arose out of* the ICISS report but has developed away from it in subsequent formulations.

<sup>117</sup> Others (Ban 2009; Hehir 2012) take the UNGA (2005) formulation as the canonical statement, but this is because these authors are concerned with the doctrine's legal status and potential for creating political change. I am less concerned with these questions than I am with the doctrine's normative foundations and details. The ICISS's more theoretically rich statement of the doctrine is more useful for answering these questions, which are completely unconsidered by the UNGA (2005). I discuss the UNGA (2005) formulation below.

<sup>118</sup> While the duties to prevent and rebuild receive 16 pages of discussion combined, the duty to react gets 32 pages.

*support* does not require *intervention*). The support duty is implied when the ICISS argues that “for prevention to succeed, strong support from the international community is often needed, and in many cases may be indispensable” and that “[i]n many cases, the state will seek to acquit its responsibility in full and active partnership with representatives of the international community” (2001, 19, 17).<sup>119</sup> While a support *duty* is not explicitly asserted by the ICISS, we shall see that it is explicitly asserted by the other canonical formulations. This suggests that we should take it to be part of the canon. Moreover, given that prevention is allegedly duty-bearers’ most important duty, support seems a crucial part of the ICISS’s R2P formulation. And asserting the existence of the support duty helps us to make sense of the obligations that are triggered when a state is not *unwilling* to protect its population, but is simply *unable*. In these cases, the most appropriate response seems to be support of the unable state in protecting its population, rather than direct protection of that state’s population from the outside.

The three core claims of R2P are quite simple. Each of the three core claims asserts the existence of one duty: the domestic duty, the intervention duty, and the support duty, respectively. The support duty and the intervention duty are what I will call the “international duties.” The domestic duty has lexical priority over the international duties, and the support duty has lexical priority over the intervention duty. This means that the support is triggered only when a state cannot or will not adequately protect its population on its own (I will explain below how this is consistent with its being a duty to help the state discharge its domestic duty), while the intervention duty is triggered only when support fails (or is reasonably expected to fail) to make a state such that it adequately protects its population. The support and intervention duties can be thought of as what Pettit and Goodin (1986, 674–5) call “godparent duties,” what Goodin (1995, ch. 5) calls “back-up responsibilities,” and what Shue (1996, 173) calls “default duties”—they are duties that arise when and only when some agent fails (or reneges, or is inadequate, or is likely soon to fail) in their own duty.

For the support duty, we should read the triggering condition of “adequate protection” quite weakly—that is, the support duty is often triggered. Thus, whenever the support duty is triggered, the domestic duty is also triggered. In this sense, it sits alongside the domestic duty. But it is posterior to it in the following sense: if the domestic duty is being *adequately* discharged, then the support duty is not triggered. In fact, depending on

---

<sup>119</sup> The “in many cases” suggests an ambiguity about whether (i) the consent of the supported state is necessary for support to be permissible (and therefore, consent is also necessary for it to be obligatory), or whether (ii) support can sometimes be permissible or obligatory in relation to an unwilling state. As we shall see, “support” may include things like conditional aid, so the second interpretation is probably the right one.

how we understand the relevant harms, support duties might always be triggered for all states—for example, if the relevant harms include the possibility of global pandemics, then all states always support all others in fulfilling these duties. The domestic duty would always be inadequately discharged if there were no support. The intervention duty, however, is triggered only in circumstances of extremely inadequate protection, namely when particularly egregious harms are occurring or imminent, and where discharging the support duty would be insufficient to stop this.

While this much is fairly straightforward, there are four important problems with the ICISS report. These relate to the relevant harms, the required actions, the duties' basis, and the international duties' bearers. First, all R2P duties are to protect people from *harms*—but which harms? This will have serious implications for the frequency and demandingness of R2P duties. Yet the ICISS vacillates. The basic principles mention “serious harm, as a result of internal war, insurgency, repression or state failure,” but this formulation occurs nowhere else in the report. Other formulations include “internal conflict and other man-kind crises putting populations at risk” (2001, XI), “situations of compelling human need” (2001, XI), or that “shock the conscience of mankind” (2001, 31, 75), “avoidable catastrophe – … mass murder and rape, … starvation” (2001, VIII), and “deadly conflict and other forms of man-made catastrophe” (2001, 19). The ICISS most commonly discusses responsibilities to prevent, react to, and rebuild after *conflict* (2001, *passim*). The ICISS clearly specifies that *military* intervention duties arise only in the face of “large scale loss of life” or “large scale ethnic cleansing,” whether “actual or apprehended” (2001, 32), but the domestic, support, and intervention duties entail much more than just military intervention.

The possibility is left open, then, that protection from a very wide array of harms is called for by R2P duties. There is an easily identifiable causal explanation for the ICISS's vagueness: it wanted to start a policy debate about action across borders to protect people from harms—*any* harms—and it used these vague yet stirring phrases because they would garner consensus, despite different agents having different interpretations of the phrases. ICISS co-chair Gareth Evans (2008, 65) claims that only “mass atrocities”—which he defines as genocide, war crimes, ethnic cleansing, or crimes against humanity—were intended to be included in the ICISS's formulation of R2P.<sup>120</sup> He notes the practical

---

<sup>120</sup> We should perhaps be sceptical of Evans' claim here. Evans' (2008) book was written after the UNGA (2005) formulation of R2P, which restricted R2P to these four crimes. It would be a convenient coincidence if, in 2001, the ICISS really meant to pick out with its vague language exactly the crimes the UNGA seized upon four years later—though it could well be true that the ICISS thought of these few crimes as amongst the bare minimum to which its principles applied.

political reasons for not trying to extend its purview: extension would water down its rhetorical efficacy and remove states' consensus on it. But we should not take states' consensus as decisive evidence that *moral* duties to protect arise only in relation to mass atrocities. The ICISS does not provide reasons or limits for the restriction, leaving the harms to which it applies unspecified.

This has left many commentators interpreting R2P as logically generating duties to protect populations from perpetrator-less harms such as natural disasters, epidemics, and chronic ongoing poverty (as surveyed in Pattison 2010, 23–4; Bellamy 2010, 150–3). This indeterminacy about the relevant harms obviously infects the content, frequency, and demandingness of R2P duties. It also infects R2P's moral basis: as we shall see, the horrors of particular harms is often taken to be a large part of R2P's moral basis—so if these particular harms are not specified, then R2P's moral basis isn't either.

A second problem arises when we ask exactly what *actions* R2P demands, particularly in the intervention duty. According to the ICISS, intervention is “action taken against a state or its leaders, without its or their consent, for purposes which are claimed to be humanitarian or protective” (2001, 8). The ICISS does not specify what makes an action “against” a state or its leaders, but perhaps the most natural reading is that such actions attempt to influence or determine a state's internal affairs, either by influencing the state's decision-making procedure or attempting to override, undermine, or circumvent that procedure.<sup>121</sup> In rough keeping with this broad definition, the ICISS also uses “intervention” to cover aid conditionalization, coercive diplomacy, political and economic sanctions, blockades, altering terms of trade, and criminal prosecutions (2001, 8).

A slippery slope opens: is any action that has influence over another agent's decision-making procedure “intervention”? If so, what is so special about intervention duties, or indeed all the R2P duties—if intervention is so broad, why should R2P relate only to particular harms, as the ambiguity around “harms” suggests? Given these two sources of ambiguity (about the harms and about what actions count as “intervention”), shouldn't the debate be a more nuanced one about what *type* or *extent* of action (whether we call it “support,” “intervention,” or anything else, and whether is it domestic or international) is warranted on the basis of *which* harms, where the harms and actions could be very minor? This would make R2P duties a very large number of graded responses to graded harms, and render a huge number of situations “R2P situations.”

---

<sup>121</sup> The focus on internal affairs reflects Article 2.7 of the UN Charter, which the ICISS (2001, 12) cites in outlining the right to non-intervention: “Nothing contained in the present Charter shall authorize the United Nations to intervene in matters which are essentially within the domestic jurisdiction of any state or shall require the Members to submit such matters to settlement under the present Charter....”

Yet such expansion is resisted in the literature and the ICISS report. As Luck stated regarding conflict prevention, if we try to “make it be all things to all people; ... in the process it could end up meaning very little to anybody” (Luck 2002, 256; similarly Bellamy 2009, 99–100; G. Evans 2008, 64–9, 71–4). And the ICISS introduced their report by claiming that it dealt with “the so-called ‘right of humanitarian intervention’” (2001, VII). Military intervention takes up more pages in the report than any other protection measure (three of eight chapters are devoted to it).

In short, the ICISS seems to want it both ways: on the one hand, they do not want R2P to relate to any and all actions in international politics; on the other hand, they want R2P to relate to much more than military intervention (e.g., aid conditionalization and so on). This again leaves unclear the content, frequency, and demandingness of R2P duties, particularly the international duties.

The third problem is that the ICISS gives an unconvincing account of the *foundations* of R2P, a convincing account of which is presumably essential for determining which agents have which duties regarding which harms (thus clearing up the first two unclarities and identifying specific international duty-bearers). The ICISS (2001, XI) lists the following “foundations” of R2P:

- A. obligations inherent in the concept of sovereignty;
- B. the responsibility of the Security Council, under Article 24 of the UN Charter, for the maintenance of international peace and security;
- C. specific legal obligations under human rights and human protection declarations, covenants and treaties, international humanitarian law and national law [the ICISS (2001, 50) specifies the Genocide Convention, Geneva Conventions, and International Criminal Court statute];
- D. the developing practice of states, regional organizations and the Security Council itself.

In general, these foundations say much more about law and convention than they do about morality. Of course, in this list the ICISS may not have been attempting to give a list of *moral* foundations. This might rather be a list of reasons why R2P arises naturally out of current laws and norms, for example. But it is informative to see just why these are insufficient as moral foundations.

Foundation (A) provides us with some guidance regarding the domestic duty: by definition, states have duties to protect their populations. But this is not the full story. As the ICISS notes, its conceptualization of sovereignty as responsibility was *revisionary* (though

not unprecedented<sup>122</sup>) (2001, 12). We want to know why we should accept this revisionism.<sup>123</sup> Elsewhere, the ICISS evokes the idea of states being “best-placed to make a positive difference” within their own borders (2001, 17). However, this is mentioned very much in passing, the notion of “best-placed” is not developed, and the ICISS does nothing to extend this idea to international duties.

Foundation (B) presumably refers only to the international R2P, since it regards the UNSC’s role in maintaining peace among states, not within them. Again, it fails to explain why, morally speaking, the UNSC has the responsibilities that the UN Charter imposes upon it.

Foundation (C) might explain why states have prudential reason to discharge the domestic duty. But the legal obligations contained in declarations of human rights and human protection cannot explain any moral duty, not least because these are largely the result of political bargaining between states pursuing their respective self-interest.

Foundation (D) is similarly unsatisfying: we want to know why it morally should be international practice to protect populations. The fact that it already is developing practice does not do this—after all, trading in slaves was once international practice and that says nothing in favour of its being morally obligatory.

Quite apart from being insufficiently explanatory, these foundations do not tell us how to determine who has the international R2P duties. Unclarity about the bearers of specific international duties is an important reason for being concerned with the doctrine’s basis, as any basis worth its salt would clarify this.

This worry applies also to the ICISS’s claim that human rights—and the broader notion of human security—are the foundation of R2P. It claims that “[t]he case for thinking of sovereignty in these [responsibility-entailing] terms is strengthened by the ever-increasing impact of international human rights norms, and the increasing impact in

---

<sup>122</sup> Francis Deng et al (1996) proposed something similar with their concept of “sovereignty as responsibility,” as did Michael Walzer, despite the latter’s opposition in general to humanitarian intervention (“a state (or government) established against the will of its own people, ruling violently, may well forfeit its right to defend itself even against a foreign invasion” (Walzer 2006, 82)).

<sup>123</sup> If we follow Goodin (1988) in holding that *all* agents have duties to protect *all* populations insofar as they can, then we can perhaps explain how a state’s duty to protect its population arises out of its claim to sovereignty. Roughly: (1) if an agent claims a right to exercise power over a population, where (2) that power would render other agents incapable of fulfilling their general duty to protect that population, then (3) that agent acquires an additional, special duty to protect that population. This requires the controversial normative premises, though, that there *is* a general duty held by all to all (rather than specific duties that arise only in the right kind of circumstances), or at least that claiming sovereignty serves to generate a special duty to protect. Allowing either of these would be sufficient for explaining the domestic duty, but neither would guide us when it comes to assigning the international duties to specific agents. So the ICISS’s foundations given above would still incomplete. Additionally, a unifying explanation of R2P would give us the same basis for domestic and international duties (rather than a sovereignty claim in the domestic case and something else in the international case)—and preferably, it would do this without having to use the contentious idea of an ever-present “general” duty of all to protect all.

international discourse of the concept of human security" (2001, 13). The ICISS describes human security as including the fulfilment of human rights, as well as people's "physical safety, their economic and social well-being, respect for their dignity and worth as human beings" (2001, 15). However, it is not clear how we are supposed to get from human rights to the existence of *duties* to fulfil those rights, particularly international duties (as O'Neill (2005) points out). Below, I will consider how we might make this move; suffice it to say here that the ICISS does not move from an unassigned "ought-to-be" claim to an assigned "ought-to-do" claim when discussing human rights.

The fourth source of ambiguity relates to the international duties' bearers. The ICISS implies that states are the key actors here, acting where appropriate through IGOs (e.g., the UNSC, International Criminal Court, International Court of Justice, the World Bank, the International Monetary Fund, and regional groups like the African Union or Arab League (2001, 22, 24, 50, 27)) or with NGOs (e.g. International Crisis Group, Amnesty International, or Human Rights Watch (2001, 21)). At times, the ICISS confers awareness-raising duties upon amorphous entities such as "religious groups, the business community, the media, and scientific, professional and educational communities" (2001, 25–6). While a proliferation of awareness-raising duties is harmless enough, awareness raising is not the main measure by which R2P duties will be fulfilled. We need to know which agents should perform—or at least lead—the *action* that is to prevent, react to, or rebuild after harms.

The ICISS gives some guidance along these lines, suggesting a "hierarchy" of responsibility for intervention duties: home state, UNSC, UNGA, regional organisations, coalitions of the willing, individual states. However, this hierarchy clearly has in mind extremely coercive forms of intervention, such as no-fly zones or military invasion. It doesn't give us general criteria that we can use to identify different duty-bearers in different contexts. Moreover, even as restricted to military interventions, the "hierarchy" is insufficiently specific. When the UNSC, UNGA, and the relevant regional organisation are deadlocked (as is so often the case), *which* coalition or individual state should we turn to? Which attributes matter? Is it pure ability to remedy the harm, or do other factors count—such as historical ties with the state in question, being welcomed by those being harmed, and so on? These kinds of general attributes can pick out particular bearers of the intervention duties (and, indeed, support duties), and explain *why* they pick out who they do.

To summarise: the ICISS is unclear about, first, the *harms* (or potential harms) that R2P duties respond to, and, second, the various *actions* these (potential) harms might warrant in various situations. These unclarities would perhaps be removed if we could

establish a convincing *moral basis* that generates all R2P duties—the domestic, support, and intervention duties. Of course, this basis would also satisfactorily explain why R2P duties exist in the first place. Finally, the basis would, ideally, enable us to assign the duties to particular agents in particular circumstances—in a way that explains why the domestic duty has lexical priority; that distributes the international duties in an intuitively plausible way; and that appropriately sets boundaries on the occurrence and demands of R2P duties.

### **7.2.2 UNGA 2005**

The second canonical formulation of R2P occurs in two paragraphs of the outcome document from the 2005 World Summit of the UNGA (UNGA 2005).<sup>124</sup> These paragraphs were “reaffirmed” by the UNSC in its 2006 resolution on the protection of civilians in armed conflict (UNSC 2006, ¶4). The UNGA and ICISS formulations agree about the three core claims: (1) the domestic duty: each state has a duty to protect its population; (2) the support duty: the “international community should, as appropriate, encourage and help” (UNGA 2005) each state in fulfilling its duty to protect its population; (3) the intervention duty: states and other international agents have a duty to, in the UNGA’s terms, “take collective action, in a timely and decisive manner, through the Security Council, in accordance with the Charter, including Chapter VII [which permits the use of force], on a case-by-case basis and in cooperation with relevant regional organizations as appropriate, should peaceful means be inadequate and national authorities are manifestly failing to protect their populations...” Again, there is a lexical ordering: the support duty

---

<sup>124</sup> The relevant paragraphs of the *World Summit Outcome* document are:

**Responsibility to protect populations from genocide, war crimes, ethnic cleansing and crimes against humanity**

138. Each individual State has the responsibility to protect its populations from genocide, war crimes, ethnic cleansing and crimes against humanity. This responsibility entails the prevention of such crimes, including their incitement, through appropriate and necessary means. We accept that responsibility and will act in accordance with it. The international community should, as appropriate, encourage and help States to exercise this responsibility and support the United Nations in establishing an early warning capability.

139. The international community, through the United Nations, also has the responsibility to use appropriate diplomatic, humanitarian and other peaceful means, in accordance with Chapters VI and VIII of the Charter, to help to protect populations from genocide, war crimes, ethnic cleansing and crimes against humanity. In this context, we are prepared to take collective action, in a timely and decisive manner, through the Security Council, in accordance with the Charter, including Chapter VII, on a case-by-case basis and in cooperation with relevant regional organizations as appropriate, should peaceful means be inadequate and national authorities are manifestly failing to protect their populations from genocide, war crimes, ethnic cleansing and crimes against humanity. We stress the need for the General Assembly to continue consideration of the responsibility to protect populations from genocide, war crimes, ethnic cleansing and crimes against humanity and its implications, bearing in mind the principles of the Charter and international law. We also intend to commit ourselves, as necessary and appropriate, to helping States build capacity to protect their populations from genocide, war crimes, ethnic cleansing and crimes against humanity and to assisting those which are under stress before crises and conflicts break out.

arises only “as appropriate” in reaction to the domestic duty (though it might very often be appropriate); and the intervention duty arises only very rarely, when the first two duties are patently inadequately fulfilled.

Yet the UNGA and ICISS formulations differ in crucial ways, rendering them inconsistent with one another. From a legal perspective, it is clear that the UNGA version has more authoritative standing (even though that standing is still weak in the context of international law as a whole (Arbour 2008)). One might therefore think that we should take the UNGA formulation as *the* canonical statement of R2P. My interest, however, is not in the legality but in the moral status and content of R2P. The UNGA leaves much to be desired in that regard, mainly because it reflects a tiny area of overlapping consensus between states. The outcome document had to be agreeable to two “tough crowds”: developing states who were concerned about infringements on their sovereignty, and developed states who were concerned about the imposition of costly duties. Both of these groups preferred to have a more minimal doctrine. This resulted in many restrictions, relative to the ICISS’s potentially wide-ranging formulation. It is evident from reports on the build-up to the 2005 Summit, and from presentations at the 2009 UNGA debate on R2P, that many of the UNGA’s restrictions are a result of states’ self-interest (Bellamy 2009, ch. 3; Ban 2009, 8; G. Evans 2008, 65–6; GCR2P 2009, 8; ICRtoP 2009, pp. 4, 7–8). It’s thus not at all clear that we should follow these restrictions for the purposes of determining the doctrine’s moral basis and moral implications.

Specifically, the UNGA made five main restrictions relative to the ICISS version. The first three helped to clarify the relevant harms and (to a lesser extent) the actions demanded from duty-bearers; while the last two further muddied the waters about the basis and conditions for international duties. First, the UNGA limited all R2P duties to protection of populations from four crimes: genocide, war crimes, ethnic cleansing, and crimes against humanity. Second, according to the UNGA, the international duty to “take collective action” (including of a military nature) arises only when a state is “manifestly” failing to protect its population from some of the four crimes. This is arguably a higher standard than the ICISS’s “unwilling or unable” (2001, 17). Third, the UNGA mentions only the UNSC (rather than regional IGOs or other organisations) as the possible authoriser of the collective action. The ICISS similarly names the UNSC as the *best current* source of intervention legitimisation, but it advocates other methods of legitimisation (e.g. UNGA or regional organisations) in cases of UNSC deadlock (2001, 53–55). The UNGA

thus circumscribes the authorisers of intervention duties to some extent.<sup>125</sup> Fourth, however, UNGA members merely say they are “prepared” to take collective action. This might imply a supererogatory voluntary intention that can be withdrawn at any time, rather than an acknowledgement of a moral requirement (as the OPGA (2009) points out). If this is the right way to understand the international duties, then this obviously has implications for the basis of those duties. Yet the UNGA does talk of the “responsibility” to “help protect” populations, rendering the source of the responsibility unclear.<sup>126</sup> Fifth and unlike the ICISS, the UNGA does not provide any criteria for permissible—let alone obligatory—military intervention, saying only that any “collective action” must be decided on a “case by case” basis.

These five differences from the ICISS demonstrate the extent of contestation around even very general formulations of the doctrine—never mind its moral basis or the method for determining the bearers of the international duties. When it comes to the basis of R2P duties, the UNGA is, in fact, completely silent—though, as noted above, there is a reading on which the intervention duty is not a duty, but a mere statement of intentions. The UNGA clearly did not see itself as spelling out the full implications of R2P, given that it emphasised the need for further consideration of the doctrine. Of course, one cannot expect much precision and philosophical insight from a two-paragraph statement that has been agreed upon by all state representatives at the UNGA. For a more nuanced approach, we might turn to the third canonical statement.

### **7.2.3 Ban 2009**

The third formulation of R2P is a 33-page report by UNSG Ban Ki-Moon, who explicitly built upon the UNGA version (Ban 2009, 4). It might be considered controversial to take Ban’s statement as R2P canon. Yet Ban’s report has been hugely influential on both the UN policy debate, and on academic output about the doctrine. For example, the Ban formulation received much support from UNGA member states at a 2009 UNGA debate on R2P (ICRtoP 2009, 4–5). Any future development of the norm—whether in an academic or policy setting—will have to take his framing of the issue seriously.

---

<sup>125</sup> Of course, this circumscription is perhaps inevitable, given that the UNGA is constrained in its proclamations by the UN Charter, according to which the UNSC is the only legitimate authoriser of force not in self-defence. Additionally, a restriction regarding authorisation does not necessarily equate to a restriction regarding duty-bearers.

<sup>126</sup> A final draft of the outcome document did state that the “international community … has the obligation to use … peaceful means … to help protect populations…” But the word “obligation” was removed after the United States ambassador to the UN objected that there is no *legal* obligation to intervene. (Hehir 2012, 79–80) Importantly for my purposes, it remains open whether the UNGA recognises a *moral* obligation. Moreover, as Arbour (2008) and Walsh and Banda (2010) argue, there *is* a legal duty, under customary international law, to protect foreign populations from some harms, namely the harms of genocide.

Ban developed “three pillars” of the doctrine, which respectively correspond to the domestic, support, and intervention duties. Pillar 1 is “the enduring responsibility of the State to protect its populations, whether nationals or not, from genocide, war crimes, ethnic cleansing and crimes against humanity, and from their incitement” (Ban 2009, 8). Pillar 2 is “the commitment of the international community to assist States in meeting those [Pillar 1] obligations” (2009, 9). Pillar 3 is “the responsibility of [UN] Member States to respond collectively in a timely and decisive manner when a State is manifestly failing to provide such protection” (2009, 9). Like the UNGA formulation, Ban’s “three pillars” have slightly different emphasis from the ICISS report. Most centrally, Ban says much less about military intervention and much more about prevention. Like the ICISS and UNGA, though, Ban’s formulation places domestic duty as prior to the international duties: the latter are there to, respectively, help the state discharge the domestic duty (presumably, only if the state actually needs help) and to perform the protective action demanded by the domestic duty only if the home state “manifestly” fails.

Like the other two formulations, Ban’s is unclear about the basis of the doctrine and about determining the bearers of the international duties. Ban claims R2P “derives from both the nature of State sovereignty and from the pre-existing and continuing legal obligations of States, not just from the recent annunciation and acceptance [in UNGA 2005] of the responsibility to protect” (Ban 2009, 9). Yet as with the ICISS’s foundations, the idea of the “nature of state sovereignty” does not go far enough, and citing legal obligations is unsatisfactory as a moral explanation (though it might well be the most feasible basis for a political consensus).

Elsewhere Ban states that “[t]he responsibility to protect, first and foremost, is a matter of State responsibility, because prevention begins at home and the protection of populations is a defining attribute of sovereignty and statehood in the twenty-first century” (2009, 10). The second half of this rationale is unsatisfactory, as the question remains: why *should* protection of populations be part of our definition of sovereignty (given that, for a long time, it wasn’t (ICISS 2001, 12–13))? Of the many ways to read “prevention begins at home,” the most natural is as saying the root causes of mass atrocities are within a home state. For this reason, the home state has the primary duty. This could be read as suggesting that this is because the home state is likely to have *caused* those root causes’ existence, but nowhere does Ban make this explicit. The best reading is probably that the most *effective* way of preventing mass atrocities is “at home”—i.e., for the home state to do so. Moreover, he refers multiple times to the importance of states having laws and policies in place that protect the “vulnerable” (2009, 12)—presumably referring not just to those who the

government might have harmed with its own actions, but those who the government is capable of protecting. All this, though, is somewhat speculative interpretation.

Ban is even more vague regarding the precise agents who bear the international duties. He seems to suggest that the main (though not only) international duty-bearers are states (2009, 29). However, he does not say anything about how we should determine *which* states have these duties in a given instance, or what role there might be for non-state actors. All in all, Ban says little more on the basis and distribution of R2P duties than does the UNGA.

All these unclarities in the three canonical formulations give us reason to cast the net elsewhere in our search for a unified, precise justification of R2P that explains the duties' lexical ordering, distributes the international duties in a clear and satisfying way, and places appropriate limits on the duties. There are some suggestions in the surrounding literature, and some potential rationales can be gathered from moral and political philosophy more generally. I now turn to these.

## 7.3 Possible Justifications and Precisifications

### 7.3.1 *Self-interest*

Although the canonical formulations leave opaque the R2P's basis and international duty distribution, a substantial literature has arisen around them, mainly within the discipline of international relations. This literature contains several suggestions for the duties' basis. One is pure self-interest: states should fulfil R2P duties because it is in their interest to do so. This might seem plausible for the domestic duty—after all, it is, in fact, in states' domestic and international interests to protect their populations.

But this is too quick. The fact that some act would advance my self-interest does not automatically imply that self-interest is why I should perform it. We should ask: if the act did not advance my self-interest, would that remove (or indeed even reduce) my reasons for performing the act? In the case of states' reasons to take measures that protect their populations, the answer is clearly "no." In fact, the historical examples that animate R2P are precisely cases where states judged it to be in their self-interest to perpetrate terrible, systematic, and far-reaching atrocities—and these states did not have any less moral reason to protect their populations. So self-interest cannot be the full explanation of the domestic duty.

What about the international duties, though? The ICISS itself introduces the duty to intervene militarily by explaining that "there are exceptional circumstances in which the

very interest that all states have in maintaining a stable international order requires them to react..." (2001, 31). In his address to the opening session of the UNGA in 1999, Annan called for a new, broader definition of national interest, on which "the collective interest is the national interest" (1999b). Evans suggests that the political will to fulfil international R2P duties might be mobilised by pointing out to states that "[e]very country has an interest in being, and being seen to be, a good international citizen" (2008, 229). James Pattison suggests mobilising such political will by drawing upon constructivist and English School theories of international relations, according to which "a state's self-interest is also determined by its identity and shared values and principles, such as the promotion of democracy, freedom, and human rights" (2010, 161).

How seriously should we take these remarks? Not very, I suggest. Pattison's comments explicitly suggest convincing states that fulfilling international R2P duties is in their self-interest in order to get them *motivated* to discharge the duties. This does not mean that self-interest *justifies* the duties. Rather, Pattison explicitly grounds R2P duties in basic human rights (2010, 23).

Evans makes explicit the distinction between *motivating* states to fulfil their international duties and *justifying* those duties. Self-interest is clearly meant to achieve only the former: his argument that fulfilling international duties is in a state's interest occurs only after he acknowledges that self-interested arguments are sometimes needed to *supplement* moral ones, though he nonetheless believes that "these kinds of straightforward moral arguments will have some resonance in every multilateral forum and every national political system." The "straightforward moral arguments" "rest ... simply on our common humanity: the impossibility of ignoring the cries of pain and distress of our fellow human beings" (2008, 229).

Annan's (1999b) address contained far more appeals to human rights and humanitarianism than to national interest. His idea that "the collective interest is the national interest" was presented as a new conception of national interest, seemingly formulated in order to *rally* states to respond to mass atrocities, not in order to justify their duties to do so.

The ICISS similarly spoke much more of human rights and human security in introducing R2P, than it did of national interest. The ICISS discusses self-interest in order to argue that self-interested motivations might move a state to act "*quite apart from* the humanitarian imperative to do so" (2001, 36, emphasis added).

Similarly, Alexander Bellamy claims that "[p]reventing future Rwandas can be boiled down to overcoming a single obstacle: how to persuade states, particularly powerful states, to risk troops to save strangers in distant lands where few strategic interests are at

stake” (2008, 429). Bellamy suggests that causing governments to develop political will requires changing the inputs they use to calculate their interests. One way in which he suggests doing this is by educating the public about the nature of overseas crises and the possibility of positive intervention. Thus Bellamy is suggesting that someone—perhaps academics or the media—*make it the case that* fulfilling international R2P duties is in governments’ domestic self-interest, by making it the case that it will earn them domestic political approval. The moral imperative to make this the case cannot itself derive from the fact that fulfilling international R2P duties is already in states’ interests, since this is the outcome Bellamy wants to achieve as an instrument to stopping gross harms.

Moreover, as Rajan Menon (2009, esp. 244) discusses at length and as suggested by Bellamy’s argument outlined above, it is simply false that it is usually in powerful states’ *overall* interests to prevent gross harms within the weaker state’s borders. This is especially the case if the state has to act alone. It is of course correct that all states have an interest in “maintaining a stable international order,” in “being, and being seen to be, a good international citizen,” and in propounding their own “values and principles.” So the quotations above have some truth to them. But these self-interested reasons are often outweighed, from a prudential perspective, by self-interested reasons *not* to protect foreign populations. (See similarly Aidan Hehir (2012, 219) on why there was an inadequate response to the crisis in Darfur.) I suggest, then, that these proponents’ arguments about national interest are best understood as having the political aim of convincing states to do what they should be doing regardless of self-interest.

### **7.3.2 Human Rights**

Along with the ICISS, many R2P proponents refer to human rights when justifying the doctrine. In Annan’s (2000) report that sparked the formation of the ICISS, the passage most commonly quoted by R2P proponents (e.g. Bellamy 2009, 35; G. Evans 2008, 31) is the one with which I started this chapter: “If humanitarian intervention is indeed an unacceptable assault on sovereignty, how should we respond to a Rwanda, to a Srebrenica—to *gross and systematic violations of human rights* that offend every precept of our common humanity?” (Annan 2000, 48, emphasis added) Ban stated that “the ultimate purpose” of R2P is “to save lives by preventing the most egregious mass violations of human rights...”; and in the UNGA’s World Summit outcome document, the two relevant paragraphs were placed in the section on “Human Rights and Rule of Law,” rather than

(say) the section on “Peace and Collective Security.” There are also many examples of human rights language among R2P’s academic proponents.<sup>127</sup>

The concept of human rights might have a large role to play in determining which *harms* R2P duty-bearers are supposed to curb. But we need a basis for the doctrine that will also explain the *duties* regarding these harms. (As we shall see, my preferred explanation is consistent with a human rights-based account of the relevant harms.) The well-known problem with human rights is that it is not clear how best to distribute their correlative duties (O’Neill 2005). Of course, there are proposals within the literature on human rights for explaining these duties’ distribution. For example, perhaps the duty-bearers of human rights are to be deduced by looking at current state practice (Beitz 2009), and perhaps this suggests that all duties to fulfil human rights are held by political entities of which states are paradigmatic instances (Barry and Southwood 2011).

This might give us an explanation for the *domestic* duty, and for why the international duties are lexically posterior to it: perhaps human rights are just the kinds of things that we claim against our home states in the first instance, and other states and international agents if and only if the home state fails. But this is precisely where we need our account of R2P’s basis to kick in: against which international agents do we then have a claim? The idea of human rights does not automatically give us an answer to this. In his “practice-dependent” account of human rights and their correlative duties, Charles Beitz (2009, 117–121) seems to suggest that the duties correlative to human rights belong to *all* states once the host state has reneged. But this solution is less precise than the one we would need to cases of, say, military intervention, where one state (or perhaps coalition) must be picked out above all others.

---

<sup>127</sup> According to Pattison (2010, 13) “the responsibility to protect is concerned with encouraging states to live up to their responsibilities to protect their citizens’ human rights” and “there is a duty to prevent, to halt, and to decrease substantial human suffering, such as that found in large-scale violations of basic human rights. This duty to prevent human suffering is not dependent on high levels of interdependence. Instead, it is universal, generated from the fundamental moral premise that human suffering ought to be tackled.” Bellamy (2009, 19, emphasis in original) claims that “sovereignty as responsibility … rests on the proposition that individuals have *inalienable* human rights that may never be rescinded. These rights are universal, not culture-specific. They are also prior to politics.” He argues that the rights and duties of sovereignty—that is, the right to non-intervention and the duties to respect other states’ sovereignty and to protect the population—can be derived from human rights (2009, 15). Jennifer Welsh and Maria Banda (2010, 218, emphasis in original) argue that R2P is based on the thought that “[t]o commit genocide or engage in ethnic cleansing is to fail to treat people *as humans*, thereby threatening humanity’s values and interests.” Evans (2008, 229) proclaims that “the basic case for R2P, and for responding in some productive way when one becomes aware of an actual or imminent mass atrocity crime, rests simply on our common humanity: the impossibility of ignoring the cries of pain and distress of our fellow human beings.” This use of “common humanity,” and Welsh and Banda’s “treat people *as humans*,” invites a human rights-based reading. Carla Bagnoli (2004, 4) simply assumes that “[a]rméd intervention is ‘humanitarian’ when it is undertaken for the sake of protecting the dignity of persons, that is, the value of their humanity. Human rights are necessary to express and exercise our humanity; they are fundamental to being a person. When we appeal to the idea of human dignity, we make a moral case for intervention, that is, one that applies universally and unconditionally.”

As I canvassed in introducing this chapter, we could say that *all* agents have duties to do what they can until someone else starts to do something, but this could result in relatively undesirable agents taking action first, or in each agent not acting on the expectation that “surely someone else will (or should).” We could instead refuse to be fussy, and insist that *any* way of distributing the duties, once it has been socially endorsed, should be followed—after all, it is better to have some distribution than none at all (Goodin 1988, 680). But given the still-emerging nature of R2P, it seems that precisely what is up for debate is which method we should use to distribute international duties to protect populations.

One option is to be pluralistic about distributing the duties. Welsh and Banda take this view of international R2P duties, claiming “there is no obvious method of determining which [duty-assigning principle] should have the greatest weight” (2010, 224; in this, they explicitly follow D. Miller (2001)). They suggest a plurality of principles: “the geographical proximity of an agent to those in need of care or protection, the special capacities of a particular agent (such as expertise or strength), and the existence of historic or special links between an agent and those in need of care or protection” (Welsh and Banda 2010, 223). They claim we should weigh up these principles in each case, using our intuitions.

Such pluralism should, I suggest, be a last resort when it comes to justifying and distributing R2P duties. Partly, this is because the justification of R2P duties will feed back into the political practice surrounding them. It would be more action-guiding for this practice if there were one clear criterion that could be consistently and reliably applied to each case. For practical political reasons, then, we should seek unity. Of course, the moral truth about R2P duties just may turn out to be complicated, pluralistic, and difficult to implement in practice. There is no reason to think that the moral truth is easy to implement. I suggest, however, that it would be not only *practically*, but also *theoretically*, more satisfying if we could identify one unified basis for the doctrine. This might not be possible, and the result might be justificatory pluralism that is difficult to implement. But before we settle for this, we should first see if we can rule out a unified explanation. I will now consider some candidate principles that might provide unifying, precisifying explanations of the doctrine’s duty distribution. To be a unifying, precisifying explanation, a principle must be a necessary, sufficient, and explanatorily satisfying condition for R2P duties, domestic and international. Particularly, we are concerned with R2P duties to protect populations, or to support others in protecting populations. (That is, we are not primarily concerned with how the *final costs* of protecting populations should be distributed, although I will mention possible upshots for this issue, when appropriate.)

### 7.3.3 Voluntary Assumption

First, it is possible that R2P duties are voluntarily incurred. Take the domestic duty: one might imagine there is an implicit contract between a state and its population, whereby that state will protect that population.<sup>128</sup> And something like voluntary assumption is suggested in the UNGA formulation of the intervention duty, in which UN member states assert that they “are prepared to take collective action.”

It is highly dubious that this is the full basis of R2P direct protection duties, whether domestic or international. We can imagine a world in which states have the powers that they have, and citizens have the dependencies that they have, but in which neither the Genocide Convention, nor the human rights declaration and covenants, nor any similar promise, had ever been made. In fact, the pre-1945 world was arguably like this. We can ask ourselves: in such a world, would states have duties to protect their populations? And do the kinds of things R2P proponents say about domestic R2P duties in the above quotations suggest that *they* would think states would have duties to protect their populations in these circumstances? Yes. Voluntary assumption is not necessary for an R2P duty.

Nor is it sufficient. If some population is undergoing ethnic cleansing at the hands of its government, such that the intervention duty is triggered, we do not think that just any state that volunteers to intervene thereby has a moral duty to do so. Imagine a volunteering state that has a long history of alliance with the rogue government, or the motive of getting access to the state’s natural resources. These kinds of reasons, fleshed out in the right kind of way, might serve to greatly outweigh any normative force that comes from voluntary assumption—even if the state was sufficiently capable of intervening successfully. More plausibly, they give the volunteering state a duty *not* to act (though they may increase the strength of its duty to compensate those who do act). Such a state might be perfectly capable of protecting the population, but even if they succeeded, we would worry about the effects on geopolitics and on the local population. So volunteering is not always sufficient to generate a duty.

Of course, one might think voluntary assumption is *sometimes* sufficient for a duty, absent the kinds of undermining reasons outlined above. This might be because voluntary assumption induces others to have certain expectations. Given these expectations, no one else will take the relevant action, and so the volunteer has a duty.<sup>129</sup> However, if the

---

<sup>128</sup> Though there are very well-known objections to this as an empirical claim, dating back to David Hume 1987[1742].

<sup>129</sup> Along these lines, Goodin (1985) subsumes voluntary assumption under the rubric of vulnerability, by viewing the *reliance* of other agents (including the person to be protected) on the volunteer as a kind of

relevant others believed that the volunteer would not be the most suitable protector (even given the volunteer's voluntary assumption), then it seems likely that these others' duties would still be up in the air. This is despite the fact that they fully expect and trust that the agent intends to take action. We could easily imagine someone else having a duty to step in to do the protecting, precisely because the first volunteer just wouldn't be very good at it. Volunteering or inducing expectations, then, is insufficient on its own for a duty.

Volunteering will, however, still have some role to play: it would be foolish to claim that volunteering never gives rise to any reason to act. For lesser harms, volunteering might *add weight* to the strength of an agent's reason to act. But it will usually be a minor supplementary reason. This reason is, I suggest, swamped by other reasons when it comes to the duties of powerful international agents to protect populations from gross harms, for the kinds of reasons just given. It is not sufficient or necessary for a duty.

### **7.3.4 Contribution**

Another possibility is contribution—that is, duties based on having contributed to the harm in question. The vast majority of R2P proponents simply do not consider a contribution-based rationale. Welsh and Banda (2010, 222) explicitly deny that contribution principle plays any role in international R2P duties, without considering any arguments in its favour. Nonetheless, it is a prominent source of duties in moral and political philosophy more generally, and could be applied to R2P. The ICISS does mention that “the rich world is deeply implicated in the [state fragmentation and state formation] process,” because arms and monetary transfers that originate in the developed world lead to civil conflicts in the developing world (2001, 5). The ICISS explains that such conflicts destabilise the political order in developing countries, increasing the risk of mass atrocities. This might be interpreted as a claim about the duties' basis, even though the ICISS does not say this explicitly. More generally, there is something highly intuitively plausible in the idea that those who have contributed a (risk of) harm should bear the cost of remedying it.

However, first, this criterion is surely not sufficient for a duty. Suppose two states are each possible bearers of an intervention duty regarding a particular ethnic cleansing. The first state has a long-standing trade relationship with the rogue government that is perpetrating the harm, such that this first state can plausibly be said to have contributed to (or at least, enabled) the efficacy of the ethnic cleansing. The second state has had no

---

*vulnerability* on the volunteer. This works partly because we imagine that no one else will act if someone has volunteered. But if no one else will act, then it seems it is not really voluntary assumption, but vulnerability, that does the work.

substantial trade or political relationship with the rogue state, but the likelihood of this second state succeeding if it attempts to intervene is very much greater than that of the first state.

It seems that any intuition about the first state “owing it” to the population in question is entirely outweighed by considerations of the second state’s efficacy. Are we so eager that the first state should attempt to repay its debt that we are willing to gamble the victims’ lives on the possibility that this attempt will fail? I would suggest not. Thus even when contribution does play some role, that role is far outweighed by considerations about how well-placed the contributor is. Contribution is, I suggest, not sufficient for an R2P duty. In fact, it is often correlated with features that make the agent wholly inappropriate as a duty-bearer: the agent that has caused a population to be unprotected is likely not to be welcomed by the local population, to have conflicting or undermining motives and intentions, and to trigger backlash and suspicion amongst other states. This goes for both the host government and international actors—so applies to all three of the R2P duties.

Second, contribution is not necessary for an R2P duty. This is particularly salient for the domestic duty: consider all those governments that intentionally and successfully operate in ways that serve to *reduce* the risk of atrocities, relative to a lack of government or relative to other policies the government might have. These governments do not cause the harm of their populations being unprotected from atrocities. Yet these governments have duties to protect their populations. The same goes for the international duty: imagine an instance of genocide, the causes of which are purely internal to the state in question, but where the home state is either unwilling or unable to do anything. Would there be no international R2P duties, because no outsider contributed to that genocide? That seems doubtful. If that’s right, contribution is not necessary for an R2P duty.

However, there is still a place for contribution in the upshots of R2P. The intuition in favour contribution seems best served in distributing the *costs* of protection, rather than the act of protection itself. There is an important distinction between having a duty to protect a population, and having a duty to pay for that protection. Perhaps those states that contributed to the harm should *fund* the protective measures, or repay those who do, for example. The duty to do this, though, is quite separate from the duty to actually do the protecting. If I break your window, I should pay for the damage—but I should not fix your window unless I am a qualified glazier. Similarly, I suggest, for collective agents’ duties to protect populations.

In sum, contribution is neither necessary nor sufficient for an R2P duty, and that the role it can play in determining protection duties is often outweighed by other moral considerations. If this is right, then it is not a unifying explanation of the duties to protect

found in the R2P doctrine. But contribution might nonetheless play a role in determining how we should distribute the final *costs* of that protection.

### 7.3.5 Association

R2P duties might instead be justified by an associative relationship between the protector and the population they are obliged to protect. Let us consider the domestic duty first. Here, the relevant associative relationship might amount to a history of mutual support between the population and the government (for example, the government's receiving taxes and providing services), or through the fact that the government's policies reflect the national culture of the population in question.

The problem with this is that governments do not have R2P duties just to their citizens, or just to those who share the national identity with which that state is most associated. Rather, they have the duties to all those in the population, that is, in the state's territory. This includes illegal immigrants, prisoners, tourists, people in transit, and so on. To claim there is an associative relationship between a government and these individuals is to stretch the concept of associative relationships almost beyond recognition.

Associative relationships are an even more problematic duty basis for international R2P duties. In this context, associative ties might include such things as similarities in religion, culture, or economic or political values, between the population to be protected and the individuals in the collective that is to protect them. Such ties might also include historical relationships between the population's state and the potential protector—such as former colonisers and their former colonies, allies in previous wars, and so on.

These ties are not necessary for an international duty: imagine a newly-discovered island, some of whose inhabitants are being wiped out by the political elite in the name of ethnic purity. The targeted population would deserve protection from some agent or other, despite a lack of associative ties with any candidate duty-bearer. However, perhaps associative ties *add to* the moral reasons an international protector has to protect a population from mass atrocities. Or perhaps, more strongly, associative ties are a *sufficient* condition for an international duty. But both of these claims go awry, at least in many cases. Many associative relationships between a population and an international agent (who is not the home state) are *historical* ties. Suppose that the current culture, political life, economy, and so on of a state has completely expunged any influence of its historical coloniser. Suppose also that the historical association gives rise to no present capacity to protect the population of that state. The associative ties rationale insists that the ex-

coloniser bears an international duty—wrongly, surely. Here the historical tie is neither sufficient, nor even contributory, to an international duty.

One could deny that *historical* association is the relevant kind of associative tie. One would then have to explain what is the relevant kind of non-historical tie—perhaps a similar current culture, or a current relationship of trust. But when we ask why *these* things matter, association itself seems to fall out of the picture. In fact, when R2P proponents argue in favour of (or elaborate on) association, their explanations emphasise associates' capacities. Welsh and Banda (2010, 224) restate the African Union's view that “‘African solutions to African problems’ have a better chance of long-term success.” Similar points are made by Evans (2008, 196), who assesses organisations of historically-related states (such as the Commonwealth of Nations and La Francophone) on the basis of their exhibited capacities. Associative ties, on their own, don't seem to be doing the justificatory work.

This does not prove that there is no way to describe the relevant kind of tie in purely associative terms, or that such ties might not somehow augment other reasons for fulfilling R2P duties. But it is highly dubious that such ties could fully explain R2P duties, since (on any plausible specification) it is unlikely that they will generate R2P duties in all cases where such duties clearly exist.

### 7.3.6 Proximity

Proximity is a potentially promising cousin of association. It is suggested frequently in the literature, when R2P proponents assert that regional actors have the primary duty once the state in question defaults (e.g., Carment and Fischer 2009, 274–5; G. Evans 2008, ch. 8; Pattison 2010, 86–87; Welsh 2009, 6–8).

As with association, however, it is important to probe this explanation further. The ICISS (2001, 54) mentions that regional actors are “more sensitive to the issues and context behind the conflict headlines, [and] more familiar with the actors and personalities involved in the conflict...” and that “neighbouring states acting within the framework of regional or sub-regional organizations are often (but not always) better placed to act than the UN...” (2001, 53). Welsh and Banda (2010, 224) argue that “closest neighbours ... are likely to respond quickly and have ‘local knowledge’ as to what kind of measures might be effective.” Similar remarks are made by Evans (2008), Pattison (2010), and several state representatives at the 2009 UNGA debate on R2P (ICRtoP 2009, 10). In short, proximity is regularly supported by outlining the *capacities* of proximate agents. But if proximity if a

mere proxy for capacity, we are better off looking at capacities directly, or treating proximity as a mere “helpful rule of thumb” for capacity.

To illustrate the inadequacy of proximity in itself as a basis for R2P duties, suppose two adjacent states are separated by a steep mountain range, with no rail links and only treacherous roads over the high mountain passes (and suppose the airplane has not yet been invented). Suppose genocide is occurring in one of the states. Suppose the other state *could* intervene—it is not ruled out by ought-implies-can—but that the second state is unlikely to succeed if it tries, and will incur a lot of cost if it tries. Would it really be the case that this state—the most proximate one—has an intervention duty, while no intervention duty is held by a more distant country that lies just over the narrow straits that can be easily navigated by ship? I would suggest not. If this is right, then proximity is neither necessary nor sufficient for an international R2P duty: the distant agent has a duty, while the proximate one does not.

## 7.4 Conclusion

The doctrine of R2P is made up of three core claims: (1) states have duties to protect their populations from gross harms; (2) other international actors have a duty to support states in discharging the duty in (1); (3) if a state is either unable or unwilling to discharge the duty in (1), then other international agents acquire duties to protect that state’s population. As we saw in §7.2, the canonical statements of these claims are insufficiently clear about the relevant harms from which populations must be protected, the required actions of international agents under claims (2) and (3), and the means by which the duties should be assigned to particular actors in particular circumstances. While the claims themselves are straightforward enough, we need an account of the claims’ justification in order for the claims to be satisfactorily precise.

In §7.3, I considered some possible justifications that would not only precisify the claims, but that would unify them under one guiding principle. These guiding principles were self-interest, human rights, voluntary assumption, contribution, association, and proximity. Each of these was fine as far as it went. But none of them went far enough. They gave us the right answers in some cases, but not in others. And the cases where they gave us the wrong answers are some of the most crucial cases for R2P—cases where no agent that can protect the population has a self-interested, voluntary-incurred, contribution-based, associative, or proximity-based reason for intervention, but where protection is nonetheless imperative. To assign particular agents particular duties in such tragic cases, we need criteria that do not depend on the chanciness of association,

causation, and so on—criteria that can be wielded in *all* cases where there is a realistic possibility of protection.

Moreover, any plausible explanation of R2P should explain the lexical ordering of the duties. Each of the bases for the duties I considered in §7.3 would be able to account for the priority of domestic duties, given that home states usually meet their preferred duty criteria better than other agents do. But these rationales have a much harder time affording *any role at all* to other states—as we saw in considering each of them, they cannot give a satisfying justification of other states’ lexically posterior duties. International duties seem to exist even when none of these five criteria are met, and these criteria do not seem to play a strong or decisive role in the precise distribution of international duties even when the criteria are met.

It remains possible that we should forgo having a *unified* basis for the doctrine. Perhaps we should evoke various of the above principles in different cases, using a pluralistic view on which none of them is necessary and none is sufficient in all cases. I have not refuted such pluralism. But I suggest that, before we consider endorsing a pluralistic justification and precisification of R2P, we should ensure that we have ruled out all *unified* methods of justifying and precisifying the doctrine. In Chapter Eight, I present my own positive proposal for such a method.

# Chapter Eight: Unifying and Explaining Responsibility to Protect

## 8.1 Introduction

In introducing Chapter Seven, I said that an adequate basis for R2P would help us answer three questions: first, why (morally speaking) the home state should have the primary duty; second, how to determine which particular agent(s) bear duties when the home state fails, and why this method of determination is the right one; and third, what the limits of these duties are. In the remainder of Chapter Seven, I suggested that the canonical formulations of the doctrine do not answer these questions, and that some possible philosophical approaches to fleshing out the doctrine do not fare well either. We were left at a loss for answering the three questions.

This chapter aims to develop a basis for RP that gives good answers to the questions. It aims to establish dependence-based duties as an precisifying, unifying explanation of R2P duties, which satisfactorily justifies the home state's normative primacy, allows us to assign international duties to agents in a clear and intuitively plausible way, and which gives us the tools to set appropriate boundaries on the demands of these duties. Along the way, it will explain how the complex twists and turns in the Dependency and Coordination Principles (as I formulated them in Chapters Two and Three) allow us to set intuitively plausible boundaries on R2P duties.

To that end, this chapter works through the various components of R2P systemically, examining each component through the lens of the Dependency or Coordination Principle (or both). §8.2 starts with the first problematic component we encountered in Chapter Seven's discussion of the canonical formulations: the relevant harms and actions. I reach back to Chapter Two's explanation of important interests and the expected value of measures regarding agents and dependents. These imply that (under a dependence-based explanation) the relevant harms and required actions are numerous and diverse. §8.3 argues that the antecedent of the Dependency Principle is necessary and sufficient for a state to bear the domestic duty, and explains how the Dependency Principle establishes (as a rule of thumb) the domestic duty's primacy. I then turn to international duties. §8.4 shows how the Dependency Principle is sufficient to generate duties in *simple* international cases—cases where one agent is best-placed, on their own, to protect a population that the agent does not govern. But not all cases are like this. In §8.5, I explain that the Coordination Principle is sufficient for producing duties in non-simple international cases. Thus, it is necessary and sufficient for an international R2P duty that

*either* the Dependency Principle's antecedent is true *or* the Coordination Principle's antecedent is true.

## 8.2 Clarifying the Harms and Actions

In outlining the R2P canon, I noted a lack of consensus about the harms R2P covers and the actions it demands. Before turning to my preferred means of justifying R2P duties, it is worth describing the harms my account covers and the actions my account potentially demands. This will allow us to understand the full potential implications of accepting my account of R2P, before I argue that we should accept this account. This will be helpful, as the subsequent argument will refer back to these implications.

Recall the view that the ICISS left possible, but did not positively endorse: that the actions demanded by R2P duties vary depending on the harm, but in principle R2P can be applied to any harms at all—not just “mass atrocities.” My account endorses a qualified version of this view, according to which R2P duties arise only in relation to sufficiently important interests. It is dubious that the only important interests are those in protection from mass atrocities. So my analysis applies more broadly than to genocide, war crimes, ethnic cleansing, and crimes against humanity—though there are excellent practical and political reasons for focusing on these in the first instance, and when it comes to publicising the doctrine to reluctant states.

Additionally, my account is not exclusively concerned with *protection from harms*, at least narrowly construed. Some of our most important interests are in the receipt of goods. Under a broad construal of “harm,” we can view the receipt of these goods as protection from the harm of not having them. Thus my account can be viewed as applying only to protection from harms. But a more natural construal of “harm” (as used by R2P proponents) is a narrower one, on which harms require a perpetrator. My account certainly applies more broadly than this.

While this contradicts parts of the R2P canon, there are good reasons to think that the canon restricts the relevant harms only for matters of “political necessity”—that is, to gain consensus between states. We need not follow the canon on this, when it comes to figuring out what the actual duties are. Even if one believes that some interests are of a different, and more important, kind than others (say, if one holds an interest-based account of human rights, and believes human rights are importantly morally different from rights of other kinds), one's list of these different and more important interests will likely include more than merely protection from the four crimes listed by the UNGA and Ban.

As stated in Chapter Two, the Dependency Principle is consistent with treating “important interests” as strictly including only some specially demarcated set of interests. “Interest-based human rights” would be a good demarcation of this set. So it would be possible, with more argumentation, to restrict the Dependency Principle’s application to R2P in this way. But such a restriction arises neither directly out of the Dependency Principle, nor (in a non-stipulative way) out of R2P. These remarks about the Dependency Principle apply equally to the Coordination Principle, which similarly employs the notion of “important interests.”

Despite this potential broadness of “important interests,” my account will not produce as many duties as it may seem. Recall that dependency and coordination duties arise only if the relevant measures would realise a positive balance of expected costs and benefits to the dependent(s), combined with expected costs and benefits to the agent(s) (where these costs and benefits might be differently weighted, according to one’s broader theory). We have important interests in a great many things, but for most of them it is highly unlikely that the Dependency and Coordination Principles will produce duties for home states or other international agents regarding them, given this “at positive value” condition.

For example, if the interest in question is “having increased funding in the public health system, so that elective surgery waiting times are shorter,” then the expected costs to international agents of taking measures regarding this interest would have to be very low for my account to produce duties for them. Most likely on my view, other states and international agents should do nothing at all, since it might be clear that the benefit to the population (shorter elective surgery waiting times) would be outweighed by the costs of friendly international relations. Here, the duty is blocked by the costs it would realise for both the agent and the dependents. And even for the home state—depending on the details of the case—there might be no duty to reduce elective surgery waiting times, if the potential funding for health could be better used to fulfil other interests of the population (say, emergency room services in hospitals). Here, the duty is blocked by the costs it would cause to the dependents.

Conversely, if the harm in question is genocide, then the duty my account produces might, on occasion, be as extreme as agents’ working alone or together in a military intervention. Although military intervention has very high expected cost (to both agent and dependent), protection from genocide is such an extremely important interest that the expected value of this measure, regarding agent and dependent, might—on occasion—render it obligatory. This liberality about required actions is consistent with the canon, which suggests a huge range of measures aimed at preventing R2P harms, of a variously

coercive and interfering nature. The difference is that my account is also liberal about the interests these actions aim to fulfil.

One last initial point relates to the meaning of “intervention.” The common sense understanding of intervention is *military* intervention, but I noted above that this was extended (in a somewhat unclear way) by the ICISS. Under the Dependency and Coordination Principles, exactly the same considerations determine the existence of both the support and the intervention duties. To determine whether there is a duty to take *any* given measure, what we need to know is: first, whether the measure is *sufficiently likely* to fulfil an important interest; second, whether its *expected value* regarding agent(s) and dependent(s) is positive; third, whether *like measures* could be taken in like cases at positive expected value regarding agent(s) and dependent(s); and fourth, whether its expected value regarding agent(s) and dependent(s) is *no lower* than any other such measure. The expected value of intervention is not necessarily going to be lower, by this metric, than the expected value of support—though of course the perceived undermining of self-determination, not to mention the loss of life involved in military intervention, are costs to be weighed in determining the expected value. These are likely to render the value of military intervention measures very low.

One might therefore think that there is no important distinction between support and intervention. This would be a mistake. The crucial difference is this. Support duties are duties to take measures either (1) to make another state such that it is best-placed to protect its population, or (2) if the other state is already best-placed to protect its population, to make it *better*-placed to protect its population (e.g. making its most efficacious measures more efficacious, or decreasing the costs of it protecting its population). Intervention duties are duties to protect a population, where the duty-bearer is not the government under which that population lives. In terms of the actual actions they demand, support and intervention duties might look very similar. But the conditions under which agents have them will look very different. The support duty is always a duty to support the state in which the population lives in discharging its lexically prior duty to protect.

Having made these clarifications about the upshots of my account, I will now argue for using the Dependency Principle to explain the domestic duty, and the Dependency and Coordination Principles to explain the international duties.

## 8.3 The Domestic Duty

### 8.3.1 The Dependency Principle as Necessary and Sufficient

To see the initial plausibility of using dependence to explain the domestic duty, simply consider the relationship we stand in to the states that govern us (regardless of whether we are decision-making members of those states in the way developed in Chapter Six). Each of us is greatly dependent on the state under whose authority we live. Our expectations, hopes, and plans are moulded by certain beliefs about what our states will do. They are clearly best-placed to fulfil certain of our important interests—like our interest in collective defence, safety in public spaces, clean air, and uncontaminated drinking water.

Of course, some of these goods are not provided by states directly, but by individuals and groups of individuals that act in a certain way at least partly because the state demands it of them. It might, then, be more accurate to say that we depend on states to fulfil our interests in the *secure* fulfilment of these important interests—we depend upon them to make it very likely that those interests will be fulfilled in a range of possible futures—and to *intervene* (through, for example, their law-enforcing members) if we find ourselves in a situation where these interests are not fulfilled.

We are, generally speaking, hugely vulnerable to our states' decisions. Given the generality and pervasiveness of individuals' dependence on the states that govern them, it would be unsurprising if this dependence played a large role in determining at least some of the duties states have to the individuals they govern.

But does the Dependency Principle give us the right answers—is it both necessary and sufficient for a domestic R2P duty? If the Dependency Principle's antecedent is *necessary* for a domestic R2P duty, then a home state lacks a duty to protect its population if the Principle's antecedent is not true of it. Can that be right? One might think not. The ICISS (2001, 13), UNGA (2005, ¶138), and Ban (2009, 8–9) all assert the domestic duty's existence *unconditionally*. If that's right, the narrow conditions of the Dependency Principle cannot be necessary for it. And intuitively, governments have duties to protect their populations even when, for example, they cannot do so for all populations that they are as well-placed to help (as the Dependency Principle's condition (4) requires).

It is true that, in the actual world, governments almost always have domestic R2P duties—that is, duties to protect their populations (even if they should sometimes be compensated for that protection by others who, for example, caused the need for protection). Governments have these protection duties so often that the duty can usually safely be stated unconditionally. But, I suggest, this is because of the nature of governments: the Dependency Principle's antecedent is almost always true of them. It

might seem that states sometimes have duties to protect their populations from mass atrocity crimes even when they are not best-placed to do so. For example, it might seem that, say, the Bangladeshi government would have a duty to protect Bangladeshis even if the US government were far more capable.

However, to believe this would be to misunderstand conditions (2) and (3) of the Dependency Principle. It would be to misunderstand condition (2) by misunderstanding the kinds of factors that can make an agent's measures more efficacious. And it would be to misunderstand conditions (3) and (4) by misunderstanding the kinds of considerations that can make it costly to take certain measures.

Regarding condition (2), measures are made efficacious not just through money, military, and international political clout. Measures are also made more efficacious if the agents taking them have experience in local institutions; knowledge of local political, cultural, and religious groups and tensions; and if the protected population trusts them (for example, if those within the population largely perceive the agent to be a legitimate ruler of the state).

Regarding conditions (3) and (4), protective measures are more costly if the agent taking them has other competing prudential demands on them. External states have strong prudential reasons to protect their own populations. It would be costly for them not to put many resources into protecting their own population, so protection of other populations is more costly for them. For external states compared to home states, it is more likely—though definitely not certain—that this cost will be disproportionate to the importance of the interests at stake.

These considerations make it unlikely either that the US would be best-placed to protect Bangladeshis, or that Bangladesh would be as well-placed to protect other populations as it is to protect Bangladeshis.

Still, what if Bangladesh were not best-placed to protect its population? Here, the Dependency Principle does give the apparently counterintuitive result that some other agent has the duty. I suggest, however, that this result appears to be counterintuitive only because, in the real world, this is very unlikely to be the case. Consider just what it would take for the dependence relation not to hold between Bangladesh and its population (that is, what it would take just for conditions (1)–(5) of the Dependency Principle to be false of Bangladesh and its population). If the Dependency Principle's antecedent is false, this means that Bangladesh is not only synchronically, but also diachronically, not best-placed

to protect its population.<sup>130</sup> It means that things would be better, regarding Bangladeshis and an outsider, if the outsider does the protecting than if Bangladesh does the protecting. It means the outsider could meet the interests not just of Bangladeshis, but of all others who are equally dependent on that outsider. Given all these assumptions (false in the actual world), we should bite the bullet and say that the outsider—and not Bangladesh—has the duty.

What about the Dependency Principle’s *sufficiency* for the domestic duty: are there any cases where, intuitively, the home state lacks a duty even though the Dependency Principle’s antecedent is true of it? If so, the Dependency Principle is not a sufficient explanation of domestic R2P duties. Given that the domestic duty is almost always, as noted above, posited as a necessary or conceptual truth (part of the “concept” of sovereignty), it is unlikely that we will find such a counterexample.

It would have to be an example where a state would bring about positive expected value, regarding itself and its population—and no less expected value than any willing agent—if it took measures to protect its population; yet it (intuitively, or according to the internal logic of R2P) lacked a duty to do so. This would be bizarre, assuming one thinks there are any domestic protection duties at all for any governments. The closest example in the vicinity is probably one where the Dependency Principle’s antecedent is true of some state, but some other state has the duty because the first state is *failing* at fulfilling the duty. But here, either the failure is due to a lack of capacity, in which case the Principle’s antecedent is false; or the failure is due to a lack of willingness or trying, in which case the home state *does* have a duty according to the Principle—but so does the second state, given the fact that we assume an agent’s own willingness when determining its duties, but are realistic about others’ willingness.

In introducing this chapter, I said that an adequate explanation of R2P would explain the lexical priority of the domestic duty over international duties. Since home states are almost always best-placed to protect their populations, a dependency-based analysis will say that the domestic duty’s lexical priority is a sensible rule of thumb. That is, it is usually true (and it might be pragmatically helpful to assume) that the home state is best-placed, and that others have duties only when the home state is unable or unwilling. This widespread truth makes the presumption of lexical priority useful for organising

---

<sup>130</sup> Of course, diachronic and synchronic duties can come apart. Bangladeshis have an important interest in being protected *now* (an urgent important interests), as well as an important interest in being protected later (a future-indexed important interests). Bangladesh might not be best-placed to fulfil Bangladeshis’ urgent important interests, while being best-placed now to fulfil their future-indexed important interests. Then, a best-placed outsider would have a dependency duty to fulfil Bangladeshis’ urgent interests, but Bangladesh would have a diachronic dependency duty to fulfil their future-indexed interests.

international politics: as a matter of practical necessary (due to time, resource, or epistemic constraints), it's usually best to assume that the home state has the duty. But the domestic duty's primacy is not strictly entailed by the Dependency Principle. The primacy is only a rough generalisation about the actual world.

One might then reason as follows: it's best to have *someone's* duty as lexically prior, and because we have assigned lexical priority to home states' duties on the basis of the rule of thumb, *all* home states now have the lexically prior duty, given that this is how those duties have been distributed. Therefore, all states—regardless of their capacities—have the lexically prior duty to protect their own populations, in actual fact and not just as a rule of thumb, because that is how lexical priority has been distributed.

This is more or less the line Goodin takes on special duties to compatriots: people's special duties to compatriots "derive wholly from the fact that they *were* appointed, and not at all from any facts about why they *were appointed*" (1988, 680). Similarly Pettit and Goodin (1986, 667) claim that "[t]here may be no rational grounds for this conventional allocation of responsibility; there may even be a duty to alter the allocation ... But in the meanwhile, the convention rules."

In affinity with my account, Goodin does assert that duties should be "assigned to agents capable of discharging them effectively" (1988, 685). Similarly, Pettit and Goodin posit a "meta-duty" to "act so that [conventional] responsibilities are optimally allocated among agents. ... This meta-duty guarantees that conventional responsibilities will be allocated (or at least should soon be reallocated) in a way that maximises the production of desirable outcomes" (1986, 673–4). But for these authors, this is a way of altering the convention, from the outside of that convention—it is not to bear directly on particular agents' duties. Rather, agents' particular duties derive from the convention—and they might do so inconsistently with the demands of the meta-duty. Nonetheless, once the convention is in play, the convention rules.

As I see it, we should only let this conventionalism take us so far. If some state is clearly not best-placed to protect its population, and it is clear who is, then why should we think that the home state has a duty to muddle along doing the best it can? Of course, we might *presume* it has the duty because it is too costly always to test which agent is best-placed. But if some state is clearly not best-placed to protect its population, we should let the particulars of the case override the lexical priority of that state's duty to protect its population. In fact, this is exactly what the lexical ordering of R2P duties is trying to assert: when a home state is (or clearly will) fail, others have duties. Contra Pettit and Goodin, I suggest that the Dependency Principle be applied directly to considerations of protecting populations—insofar as our epistemic and resource constraints allow.

And these constraints might not allow much. In practice, we might only know that a state is not best-placed, or is unwilling, when a horrific atrocity is already unfolding. In other cases, we will know a state is very poorly placed indeed, but all other agents will be even more poorly placed. (In these cases, a partial reason for the other agents' being very poorly placed might precisely be that the home state's inability or unwillingness would impede any potential actions of outsiders—North Korea might be a case in point.) Other times, an outsider will clearly be best-placed, but the measures it should take, given costs to itself and the population at risk, will fall well short of military invasion. For these reasons, rejecting the absolute primacy of the domestic duty does not automatically open the floodgates to extensive international action. Its practical implications, then, are likely to be endorsed by those who agree with Pettit and Goodin that we should hold more tightly to the lexical ordering.

### ***8.3.2 Complications with Domestic Duties***

My detailed formulation of the Dependency Principle provides a number of tools for addressing complex instances of domestic duties. (I will address international duties shortly.)

First, consider the diachronic capacities of home states, alluded to above in the Bangladesh example. Both Ban and the ICISS discuss the importance of governments' duties to build their capacities (ICISS, 22–27; Ban 2009, 10, 27–8). The UNGA (2005, ¶139) posits such duties explicitly. My account provides a good explanation of these duties. It views them as instances of diachronic dependence: the government is now best-placed to protect its population at some future point, taking into account the long-term costs and benefits, to itself and its population, of it making itself synchronically best-placed to do so. (Of course this does not necessarily apply only to the home state, but to any agent that could, diachronically, make itself best-placed to protect this population.)

If the government is not now best-placed to make itself synchronically best-placed in the future, so that the expected value of it protecting its population at that future time is lower than that of some other agent, then it would, I believe, be very doubtful that it had a duty according to R2P's proponents. Probably in all real-world cases, such a government would be a failed state or near-failed state, and would lack a duty for that reason. (As Erskine (2001) points out, failed states lack agency altogether, so cannot have any duties.)

Second, domestic R2P duties are a good example of “packaging” interests. The Dependency Principle is primarily designed for cases where there is just one person whose interests are up for fulfilment, but it importantly allows otherwise. According to the

Principle, if two measures are equally most likely to fulfil an important interest, but one of them also fulfils *other* important interests, then the latter measure is the “most efficacious” measure for fulfilling that interest.

To demonstrate this in Chapter Two, I gave the stylised example of Peter and the 100 drowning toddlers. It might be far too costly for Peter to jump in after each one (he will die of hyperthermia, say, which is too much for morality to demand). But he could build a fence around the pond. Suppose the fence would not realise positive expected value, regarding agent and dependent, if it would save only one child (say, because the construction process is extremely dangerous for Peter). But if it would save 100, its net benefits for dependents would outweigh its net costs to the agent. Clearly, Peter should build the fence if there are 100 children, but jump in after each child if there are only a few children. (In-between cases will require difficult value judgments, and it might be indeterminate how many children’s lives the dangerous fence construction is worth.)

So it is with the state: if only one person’s interests were up for fulfilment, then the complex system of national defence, criminal and civil law, and so on that protects important interests would be too costly. Likewise, protecting each person’s interests on some sort of piecemeal basis would be too costly, given how many people there are. By packaging interests, though, we get the right result: states can, and should, protect everyone in their populations, taking non-piecemeal measures.

Third, other agents’ duties might affect domestic duties. For example, suppose the home state is best-placed to alleviate some harm, but an external agent entirely caused the harm. The external agent might have a contribution-based duty to incur (some of) the costs of remedying the harm, say by financially compensating the state. If it is likely that the external agent will be made to compensate the state—say, by order of the International Criminal Court—then this would reduce the costs to the home state of alleviating the harm. This will likely ensure that the home state would produce positive expected value for itself and the population by alleviating the harm, so as to be eligible for a dependency duty to do alleviate it. This is all included in the lights of the Dependency Principle.

## 8.4 International Duties: Single-agent Cases

### 8.4.1 Initial Evidence for a Dependence-like Ground

While the domestic duty is usually asserted unconditionally in the canon and the commentary, the international duties are often accompanied by hints about which particular agents might discharge those duties. These hints often evoke the dependence-like

concepts of capacity, ability, or opportunity. In discussing who might support home states or intervene when they fail, Ban talks about agents that are “well placed” and “capable” (2009, 15, 31). Evans states that “it is the *ability* and *willingness* of national governments to act as good friends and neighbors, as generous donors, as persuasive diplomats, … as appliers of coercive pressure and, in really extreme cases, military intervention, that is crucial to any solution [to ‘catastrophic human rights violations’]” (2008, 196, emphasis added). The ICISS predicts that “[w]hat will be increasingly needed in the future are partnerships of the able, the willing and the well-intended – and the duly authorized” (2001, 52). The ICISS argues that the UNSG has a particularly strong role to play in discharging international R2P duties, because this role-bearer’s authority and fame provide him a “unique opportunity to mobilize international support” (2001, 72–3).

However, as can be gleaned from the above quotations, R2P proponents persistently speak of what *might actually* be done in the same breath as they speak of what *should* be done. In other words, R2P proponents often switch between a predictive mode (e.g., when mentioning agents’ willingness) and a normative mode (e.g., when mentioning agents’ capacities or authorisation). The conflation of these modes almost certainly explains why these authors mention both willingness and capacity as being important: these are two necessary conditions for action *actually* being taken in the future.

But when considering what action *should* be taken by some agent, we can put aside the question of that agent’s willingness: this is relevant for the predictive mode, but not the normative one (though of course the unwillingness of one agent has implications for the duties of *other* agents). The same goes for authorisation—authorisation is, of course, politically important. But our question here is the logically prior one of which agents are worthy of receiving authorisation—and, even more than this, which agents should have not merely an authorisation (i.e., prerogative) to act, but a duty (i.e., requirement) to do so.

Once willingness and actual authorisation are put aside, we are left with capacity as a possible source of duties. Of course, it is possible that R2P proponents mention capacity only in the predictive mode. After all, we should not believe that someone *will* act if do we not believe they are able to act—so capacity certainly bears upon the predictive question. However, my suggestion here is simply that capacity’s persistent appearance in the literature gives us good reason to *consider* whether it would work as a moral justification of the international duties, in a way that willingness and authorisation do not seem to work.

### 8.4.2 The Dependency Principle as Sufficient

To start, consider simply a case in which one agent (whether state, IGO, NGO, or so on) is best-placed to protect some population within a state, where the best-placed agent is not that population's home state.<sup>131</sup> Assume it would not be better, regarding agents and dependents, if this best-placed agent acted in concert with others, or supported the home state in protecting the population. Call this the "simple intervention" case.

The Dependency Principle asserts a duty in all simple intervention cases. But is being best-placed sufficient for an intervention duty in the simple case? I suggest so. For the best-placed agent to have a dependency duty in this case, its protection would have to be more morally valuable for population and protector than the home state's—even diachronically, and even taking into account the population's probable interest in being protected by their home state. There are two things to note about cases with this structure. First, they are very rare in the actual world. That the Dependency Principle posits a duty in all simple intervention cases does not mean that it is constantly condoning infringements of states' sovereignty, because (among other things) the measure must be proportionate to the importance of the interest and must produce positive expected value regarding the protector and the population. Second, it is exactly these uncommon cases that R2P was originally designed to address. If there were no duty in a simple international case, the R2P would be almost entirely inert. So if our intuitions deny there is a duty in such a case, our intuitions deny R2P. That would be no argument against the Dependency Principle being a good explanation of R2P.

We can consider the same kind of case, but where the outside agent's duty is a *support* duty rather than an *intervention* duty. Now, one might think that the Dependency Principle on its own cannot produce support duties, because support implies cooperation, and the Principle explicitly only deals with measures that are likely to fulfil interests *if no other agents cooperate*. But it all depends on how we specify the relevant interest.

If we take the population's interest not in "protection" but in "being governed by a state that is best-placed (or better-placed) to protect them," perhaps an outside agent is best-placed to meet this interest on its own. If the home state is not best-placed, on its own, to make itself such that it is best-placed (or better-placed) to protect its population,

---

<sup>131</sup> I will speak of "populations" because it seems very much more likely that the aggregated interests of many individuals (or the population's interests taken as a whole, if it's possible to see the population's interests as more than the sum of the individuals' interests) will be sufficiently important to render international action obligatory. I do not mean to deny the possibility of international agents having dependency duties to take measures to fulfil the interests of just one individual. But these are not the paradigm case for R2P, nor the most common, likely, or pressing example of the Dependency Principle applied internationally. I will also simplify matters by assuming that the population lives under just one state, rather than spanning several states.

then it cannot have a duty to make itself best-placed (or better-placed) to do so. But an outside agent might be best-placed to make the home state best- or better-placed, even if the home state doesn't cooperate. Call this the "simple support case." It is one where one agent is best-placed, on their own, to make another agent best- or better-placed to protect its own population.

The outsider would then—on the basis of the Dependency Principle—have a support duty: a duty to make another agent such that it is best-placed (or, if already best-placed, better-placed) to protect its own population. In Chapter Two, I gave a similar example: if it is important for me that my parents can love me, then other agents might have dependency duties, based on that agent-indexed interest, to make it the case that my parents are capable of loving me, or to make them more capable of loving me. The support duties of R2P—at least in simple cases—work analogously. Again, this hypothetical is highly idealised: most often, agents are best-placed to change themselves, or, at the very least, it is best that agents *coordinate* with other agents in the pursuit of self-change. I take it as a central claim of R2P that, in the simple support case, the outside agent would have a duty. The Dependency Principle can capture this thought.

Here I have suggested only that the Dependency Principle is *sufficient* for generating duties in the simple intervention case and simple support case—that nothing in addition to the Principle's antecedent needs to be true in these case in order for there to be a duty. I have not considered the question of whether it is necessary. As we shall see (and as the reader may already have seen), the Dependency Principle is actually not necessary for all international R2P duties. This is because most of these duties are not simple, single-agent cases. So I will not consider whether the Dependency Principle is necessary. Yet before addressing the issue of multi-agent cases, it will be useful to say how the Dependency Principle's treatment of the simple, single-agent cases interacts with another, better-known treatment of them.

#### **8.4.3 The ICISS's Criteria for Intervention**

It is common for R2P proponents—and philosophers of war more generally—to produce lists of criteria for when intervention (particularly military intervention) is permissible or obligatory.<sup>132</sup> The ICISS, for example, spends several pages developing its list of criteria for

---

<sup>132</sup> The High-Level Panel on Threats, Challenges, and Change (2004) lists: seriousness of threat, proper purpose, last resort, proportional means, and balance of consequences. Pattison (2010, 109) lists: discrimination between permissible and impermissible targets, proportionality, the non-use of certain kinds of combatants (e.g. child soldiers), and fulfilling a duty of care of one's own soldiers. Walzer (2006, 90) lists the "just cause" criteria of an obvious lack of "community of self-determination", such as "in cases of

when military intervention is “legitimate”: just cause, right intention, last resort, proportionality, and reasonable prospects (2001, 32–7). These criteria are designed for cases analogous to the simple intervention case described above, so it will be useful to discuss these criteria as a contrast to the Dependency Principle’s treatment of these cases.

Before doing so, however, I should note it is disputable whether any adequate explanation of R2P would be absolutely required to align with the judgements produced by the ICISS’s criteria, since the other documents in the R2P canon (UNGA 2005; Ban 2009) do not say anything about specific criteria for intervention. The UNGA recommended a “case by case” approach to determining whether intervention (military or otherwise) is required, rather than setting a general “trigger” or “threshold” for the loss of any rights against intervention (military or otherwise). Ban (2009, 22) similarly argues that “there is no room for a rigidly sequenced strategy or for tightly defined ‘triggers’ for [military] action.” And in the ICISS’s consultations, many were not convinced that such lists of criteria for permissible military intervention were a good idea (Bellamy 2009, 45–6). For these reasons, we should probably be sceptical of general lists of criteria for when military intervention is permissible or required—and sceptical of whether a moral basis for R2P would have to align with the ICISS’s list.

Moreover, it would be very strange if *military* intervention were legitimate exactly when certain criteria were met, while other types of intervention and support had to be determined on a case-by-case basis (as is suggested by the lack of ICISS criteria relating to these other measures). And even if the criteria were the necessary and jointly sufficient conditions for justified military intervention, there would still be huge problems in knowing whether the Dependency Principle’s substantive judgments aligned with theirs, since their phrasing is sometimes vague.

Even given these reservations, one might reasonably wonder how the Dependency Principle fits in with the kinds of criteria that are generally thought necessary to render military intervention permissible or obligatory. And one might think the fit is poor, since the Dependency Principle gives a formula for determining when certain measures are *obligatory*, but has nothing to say about whether intervention (or support, or indeed protection measures of any sort) might be merely permissible.<sup>133</sup> The ICISS introduces five

---

enslavement or massacre.” McMahan (1996) similarly addresses the just cause issue, claiming that a just cause is a situation of severe human rights violations where the persecuted group welcome the intervention. (For Walzer and McMahan, the implication is that their more general principle of justice before, during, and after war apply to military interventions once the intervention begins.)

<sup>133</sup> Though the distinction between permissible and obligatory intervention might have no practical importance: some theorists assert that intervention is permissible *only if* it is obligatory, at least if the intervention is military in nature (Rodin 2006; Tan 2005). The thought goes that intervention generally has

criteria for “legitimate” military intervention, which presumably amounts to permissible, but not necessarily obligatory, intervention. (However, the ICISS certainly thinks there are *sometimes* obligations to intervene, and the “legitimacy” criteria would have to be met by any such obligation-bearer.) Additionally, the Principle can be applied to all sorts of protection measures, not just ones involving military intervention. Thus the Principle and the criteria seem sharply at odds.

In fact, the dissonance between the Dependency Principle and the ICISS’s criteria—although present—is much less extensive than it might appear. If we endorse the Dependency Principle as the basis of R2P, the ICISS’s criteria actually become a good check-list. Two kinds of case are particularly good at illustrating this. The first are cases where the atrocities are unfolding rapidly, and agents must quickly figure out whether military intervention is required and, if so, by whom. The emphasis is on speed. The second are cases where several agents must agree on whether intervention is required and, if so, by whom (these may or may not also be speed-requiring cases). Here, criteria can serve as “salient points” to guide deliberation. The ICISS’s criteria are a useful *guide* to quick decisions, and to multilateral deliberation, because it will almost always be the case that, if the Dependency Principle’s antecedent is true, then something like the ICISS’s criteria for military intervention will also be true, and vice versa. (However, the “vice versa” claim is not necessarily true if the ICISS’s criteria relate to merely permissible, rather than obligatory, intervention—though see fn. 133).

Of course, checking off the criteria would require some difficult deliberations of just the kind involved in directly determining whether the Dependency Principle’s antecedent is true. But the criteria might nonetheless helpfully guide agents through these deliberations, by pointing towards some typically relevant considerations. Thus, by viewing the criteria through the lens of the Dependency Principle, we can see the grain of truth in them. This grain of truth is that they are useful rules of thumb for military intervention. Given the kinds of reservations outlined above, we should be wary of taking the criteria as anything more than rules of thumb. But the Dependency Principle gives them to us in that capacity. To see how the Dependency Principle gives us these rules of thumb, let us consider each in turn.

---

such high costs—for everyone involved—that it could only be permissible if it is to prevent a very terrible harm. But then, if the harm is as terrible as all that, then intervention must be obligatory. Another way of putting the thought is this: how could things ever be so bad that intervention is permissible, without things being so bad that intervention is also obligatory? How could intervention, in such a terrible case, be morally optional? If this is right, then the criteria given by, for example, the ICISS can be taken as criteria for obligatory, not merely permissible, intervention.

First we have “just cause.” In the ICISS’s view, for military intervention this is limited to: “large scale loss of life, actual or apprehended, with genocidal intent or not, which is the product either of deliberate state action, or state neglect or inability to act, or a failed state situation; or … large scale ‘ethnic cleansing,’ actual or apprehended, whether carried out by killing, force expulsion, acts of terror or rape” (2001, 32). For the Dependency Principle’s antecedent to be true, the measures in question must produce positive expected value for protector and population. If the measures are military intervention, this will probably be true only if the interests that the measures will fulfil are extremely important—as protection from the ICISS’s two listed harms is.

Of course, arguably protection from other harms, such as large scale loss of life from a natural disaster, is *as* important as protection from these two harms. But in cases of natural disaster, it is unlikely to be true that military intervention is the *most efficacious* measure available—and the Dependency Principle deals only with the most efficacious measure.<sup>134</sup> Nonetheless, it remains possible, by Dependency Principle lights, that military intervention will be required for harms other than the two the ICISS lists. So the existence of one of these two harms should, by the Principle’s lights, be used only as a presumptive necessary criterion for legitimate military action.

Second, “the primary purpose of the intervention must be to halt or avert human suffering” (2001, 35). It is possible that the ICISS was concerned with intentions for their own sake, but it seems far more likely that they were concerned with the *effects* of acting on various intentions. The duties produced by the Dependency Principle are duties to intentionally take measures, where the measures are likely to fulfil important interests and the measures would realise positive expected value regarding the relevant parties. This seems to accord with the ICISS’s criteria. However, it is possible to intentionally take certain measures, where those measures are likely to fulfil important interests, without having the intention of fulfilling those interests. For example, I can intentionally put out your house-fire in order to protect my rug (which you are borrowing) from being burned. In the process, I might save your life. I have intentionally taken an action which has saved your life, but that was not my “primary purpose.”

This might look like a deep source of dissonance between my account and the ICISS’s. Yet while dependency duties are just duties intentionally to take measures, the measures are selected precisely for their likelihood of fulfilling important interests and their positive expected value—that is, for their effects. The measures required to fulfil important

---

<sup>134</sup> By “military intervention,” I mean “violent military intervention.” That is, I mean to exclude cases where the military is used to deliver aid after a natural disaster, without the intention or expectation of using violence.

interests will almost certainly be incompatible with the measures required to act on the intentions the ICISS wants to rule out—intentions such as altering borders, advancing a particular group’s self-determination claim, overthrowing a regime, or occupying a territory. These purposes will all require different—subtly different, but different nonetheless—measures to those that the Principle will pick out to fulfil important interests. And on the occasion when the measures for acting on those intentions really do coincide with those prescribed by the Principle, it would be very odd to say that the measure is thereby prohibited. After all, the costs to the population of (for example) overthrowing the local government have already been taken into account in the principle’s prescription of the duty.<sup>135</sup>

The third ICISS criteria is “last resort”—“there must be reasonable ground for believing that, in all the circumstances, if the [diplomatic and non-military] measure had been attempted it would not have succeeded” (2001, 36). This accords with the Dependency Principle, according to which the duty-demanded measure must be the most efficacious measure open to the agent to fulfil the interest at stake. Given the high risk of long-term and great destruction—not to mention the high casualty rate—that comes with military intervention, it will almost certainly only be the most efficacious means if all other measures would not be efficacious at all.

Fourth, the ICISS lists “proportional means”: the military measure should be “the minimum necessary to secure the humanitarian objective in question. The means have to be commensurate with the ends, and in line with the magnitude of the original provocation” (2001, 37). The “minimum necessary” idea is captured by the Dependency Principle’s “most efficacious means” criteria, as I just described it. The ICISS’s idea that the means must be “commensurate with the ends” is handled by conditions (3) and (4) of the Dependency Principle. Applied to military intervention, these would say that military intervention must realise positive expected value regarding protector and population in this case, *and* that there would be positive aggregate expected value if the agent acted on all relevantly similar cases.

Finally, we have “reasonable prospects”: “[m]ilitary action can only be justified if it stands a reasonable chance of success...” (2001, 37). Again, this is directly addressed in the

---

<sup>135</sup> I should be fully explicit about the bullet I’m biting here. Suppose State A can invade State B, thereby preventing genocide and restoring order; but that intervention will have the effect of restoring the oil pipeline that provides State A’s principal source of fuel. Suppose that, when we weigh up the value and likelihood of these two effects, we get the result that State A should take the measure. According to my account, State A does their duty if they intentionally takes the measure—even if, in taking the measure, they only intended the second, and not the first, effect.

Dependency Principle, where condition (2) states that the measures are likely to full the interest in question with a likelihood that is *proportionate* to the importance.

## 8.5 International Duties: Multi-agent Cases

### 8.5.1 Calls for Coordination in the Literature

In using dependence to justify, precisify, and unify R2P, I have so far talked about domestic cases and simple international cases (i.e., international cases in which one agent is best-placed). This is hugely unrealistic: in the vast majority of real-life international cases, it will be best if a number of states work together. Not only will it be *best* if agents work together: it will only be *sufficiently good* (i.e. realise positive expected value regarding relevant parties) to generate a duty if agents work together. This is especially so for the support duties: usually, positive expected value will require mutual responsiveness between supporter and supportee, as well as between multiple supporters. The Dependency Principle does not produce duties in these cases. It also has trouble with non-simple intervention duties: what if each of two states would realise negative expected value if each intervened to protect the state's population, but by doing so together they would produce positive expected value? Or what if the positive value would be higher if two or more agents worked together in intervention? The Dependency Principle cannot answer these questions. Because of this, the Dependency Principle's antecedent is not *necessary* for an international R2P duty, though it is sufficient to explain the domestic duty and simple international duties. We need another set of sufficient conditions to capture complex international duties—those in which agents must work together.

Indeed, the R2P canon makes very clear the importance of various kinds of multilateralism. The ICISS claims that “the issue of international intervention for human protection purposes is a “clear and compelling example of *concerted action* urgently being needed,” that “[i]n key respects … the mandates and capacity of *international institutions* have not kept pace with international needs or modern expectations” (2001, 3, emphasis added); and that coordination is a “perennial concern” (2001, 3, 26). These coordination duties appear to be held by a wide range of agents, including NGOs, who should “learn how better to *coordinate* among themselves, *mobilize* constituents globally, *work with* the media, and *move* governments” (2001, 21, emphasis added).

Of course, one can make too much of this language: the ICISS was eager to foster consensus among states, and platitudes about multilateralism increase the likelihood of such consensus. But there is good reason to think that, at least to some extent, the ICISS

meant what it said. After all, it made some concrete suggestions for possible ways of coordinating. It suggested that the UNSG should assign new roles within the UN (thus altering the decision-making procedure of a collective of which he is a member), by asking UN Member States to give “regular reports and updates on capacities, capabilities and current practices designed to prevent conflict...” The ICISS also suggested that a new collective agent be created—an “integrated Task Force”—to assist and recognise conflict prevention efforts by those states that are vulnerable to conflict (2001, 26, 27).

The language of coordination also appears in other parts of the canon. The UNGA states a preparedness “to take *collective* action, ... through the Security Council, ... in *cooperation* with relevant regional organizations as appropriate” (2005, ¶139, emphasis added). Ban expresses the need for collective agents to consider ways to coordinate (e.g., providing “collective international military assistance” (2009, 18)), reform supra-collectives of which they are members (e.g., increasing the rule of law assistance that the UN offers its members (2009, 21) or sharpening the focus of the Human Rights Council (2009, 11)), and form new collective agents (such as a “standing or standby rapid-response civilian and police capacity” (2009, 18)).

Such wide-ranging calls for responsive action occur equally in R2P commentary. At the 2009 UNGA debate on R2P, state representatives generally agreed that the international duty was a “collective” one (GCR2P 2009, 2, 6-7; ICRtoP 2009, 6, 9–10). Several states that supported R2P also supported the formation of rapid-reaction forces, neutral arbiters, or similar (Hehir 2012, 247). The High-Level Panel on Threats, Challenges, and Change (HLPTCC), in its report that built on (*inter alia*) ICISS’s findings, suggested that the UNSC install a system of “indicative voting,” whereby members would have to publicly declare their positions on a resolution before the vote on that resolution occurs, and that use of the veto of the five permanent UNSC members “be limited to matters where vital interests are genuinely at stake” (HLPTCC 2004, 82)—again, calling on agents to reform a collective of which they are a member.

Similar suggestions come from academic proponents, who have less reason for political pandering. Bellamy advocates the formation of a permanent “embargoes unit” to systematically guide the imposition of targeted sanctions (2009, 141–6); the centralisation of early warning about mass atrocities in the UN Department of Political Affairs (2009, 109–10); and increasing the capacities and jurisdiction of the International Criminal Court and the Human Rights Council (2009, 127–8). Evans encourages states to “join and participate actively in international organizations and regimes” (2008, 89); sign peace accords (2008, 110); and change their own collective decision-making procedures in numerous ways to better prevent mass atrocities (2008, ch. 4). Pattison suggests enhancing

the United Nations Stand-by Arrangements System, which allows states to commit to contribute troops to UN operations (2010, 227–9); creating a cosmopolitan UN force with corresponding cosmopolitan democratic institutions (2010, 233); and broadening the legal powers of regional organisations (2010, 65–6, 236–9).

### **8.5.2 The UN as Sole Duty-bearer?**

Despite all this, one might think that we need not posit duties for agents to work together, since the UN is best-placed to fulfil all international R2P duties, at least in a great many cases.<sup>136</sup> At the very least, the UN might have a duty to *change itself so that* it can bear international R2P duties. In Chapter Three, I described the conditions necessary for a collective to have a duty to change itself: the change must be consistent with its current goals, and its current decision-making procedure must be such that constituents can use their roles within the collective to operate the procedure in such a way that the decision-making procedure changes itself. The UN does have human protection as a current goal (ICISS 2001, 13), and regularly deploys military personnel for peacekeeping purposes (though it relies on member states for troops) (United Nations 2009). And it is possible that members could work “from the inside,” as constituents, to change the UN’s decision-making procedure so that it can bear international R2P duties in all instances, including, for example, by establishing a UN standing army.

Yet such a change to the UN’s procedure would be quite radical. Take, for example, the method by which the UN authorises the more extreme forms of intervention (e.g. no-fly zone imposition, military invasion). All decisions over these measures lie with the UNSC. And all resolutions of the UNSC are subject to veto by the five permanent members (P5), who regularly block intervention proposals on the basis of their own political ends. A UNSC duty to change its procedure such that it can fulfil all R2P duties would mean that each member has a duty to use UNSC deliberative fora to introduce and approve a proposal whereby, say, the UNSC has a rule that the P5 veto may not be used when the UNSC is passing resolutions that relate to international R2P duties and the potential vetoing party does not have a vital national interest at stake. (The ICISS (2001, 51) advocates such an agreement.) Any change in the P5’s powers would require the consent of each of the P5. This is a deep change to the UN power structure, which is very unlikely to be forthcoming.

---

<sup>136</sup> The ICISS even states that “[t]he world already has in a place a standing military and diplomatic organization with the capacity (if not always the will) to deal with the whole spectrum of peace, security and human protection issues: we call it the United Nations.” (2001, 48)

Nonetheless, on my account, this does not get the UNSC or P5 off the hook: members should still do what they can to make the UN such that it can fulfil its diachronic international R2P duties. R2P proponents almost unanimously agree (Hehir 2012, 229–54; Pattison 2010, 54–7; Bellamy 2009, 75–6; GCR2P 2009, 2; HLPTCC 2004, 65).<sup>137</sup> Yet the fact that the UN is very *unlikely* to adequately reform itself in the near future should be taken as a given when we are identifying the duties of *other* agents. How can victims of mass atrocities—or anyone else—be satisfied with the prospect that if no agent alone is both sufficiently well-placed and willing to protect them, but agents together are, then no one has a duty regarding their protection? Or that agents have no duties to work together if doing so would protect some population—or indeed *all* populations—more fully? This would be morally unacceptable.

### **8.5.3 The Coordination Principle as Sufficient**

The Coordination Principle deals with duties in a slightly different way to the Dependency Principle. While the latter starts with “important interests” and the best placed agent, the former starts with a particular state-of-affairs, “p,” and a best-placed *set* of agents. The “p” can be any non-actual state-of-affairs in which important interest(s) is (are) fulfilled. With this in mind, make p “that population P is adequately protected from harm H,” where each member of the population has an important interest in being adequately protected from H.<sup>138</sup> A population’s being “adequately protected” means that the likelihood of H in that population over some timeframe is below some threshold. This threshold will most likely be set by context—crucially, the size of P, the badness of H, and whether it is possible that (if agents try) the likelihood of H in P will be reduced (that is, “adequate protection” should be, in some broad sense, feasible).

Regarding the adequate protection of particular populations, where that protection would require states to work together, the Coordination Principle is sufficient to generate the multitude of international duties that R2P proponents advocate. According to the

---

<sup>137</sup> Some disagree. Welsh and Banda express scepticism about the efficacy of assigning *any* international R2P duties to the UN, claiming that “by assigning the collective responsibility to the UN, responsibility is less likely to be distributed down to where it needs to be felt and exercised: by individual states.” (2010, 225) This is not such a large worry if we accept my reductively individualist account of collective duties, on which duties of collectives are duties of individuals to use the decision-making procedure to distribute certain roles, and then to perform those roles. Thus we can, and should, explain UN failure entirely in terms of state failure. Of course, states might try to resist this, in just the way any moral agent might try to sweep their moral failings under the carpet. But responsibility for failure to act would not get “lost” at the UN level: each member state has a duty to join peace-keeping missions, until enough others do. If not enough do, all who have not bear responsibility.

<sup>138</sup> It might be a single individual rather than a whole population, but the latter example is more morally pressing and in keeping with the *mass* harms with which R2P proponents are generally concerned.

principle, a wide variety of individual and collective actors—states, NGOs, UN ambassadors, individual political commentators, religious leaders, and so on—have a duty to act responsively to others, either with a view to adequately protecting given populations or with a view to there being a collective that has a perfect duty to adequately protect given populations. If unilateral responsiveness would be pointless, they have a duty to signal conditional willingness (i.e. willingness to be responsive if others signal willingness to be responsive), and they have duties to be responsive if others signal likewise.

These duties are reducible to duties of individuals acting within a group context. To see this, consider the series of claims that can be made by individuals in inadequately protected populations, if the Coordination Principle's antecedent is true. First, in cases where unilateral responsiveness would be pointless, these individuals each have a claim upon *each* bearer of a coordination duty that they signal conditional willingness. Second, if and when enough duty-bearers have signalled conditional willingness, such that responsiveness would not be pointless, individuals in unprotected populations each have a claim that duty-bearers take responsive steps. Third, if responsiveness results in a collective (or aggregate of responsiveness individuals) that is capable of adequately protecting them, then they each have a claim that the collective (or each member of the aggregate) takes steps to protect them. In cases where a collective is formed, individuals in unprotected populations each have a claim that each constituent of the collective uses his role in the collective with a view to the unprotected population being adequately protected. If responsiveness results in an aggregate of responsive individuals, then individuals in unprotected populations each have a claim that each member of the aggregate is responsive with a view to the unprotected population being adequately protected.

The Coordination Principle is an extremely versatile tool for dealing with a wide array of p values. We can, for example, entertain states of affairs in which *all* populations are adequately protected. These are states of affairs might be brought about through such steps as “P5 UNSC members agreeing not to apply their veto to resolutions regarding mass atrocities in cases where there would otherwise be a clear UNSC majority”; or “There being an independent judiciary, formed by the UNGA, for assessing the appropriate reaction to mass atrocities in cases where the UNSC is deadlocked”; or “The UNGA agreeing to form a standing UN army, that can be mobilised by the UNGA (or the UNSC, or...), to which individual soldiers can sign up directly, and that is funded by UN member states.” The possibilities are numerous. For each of these reforms—and I lack space to go into the details of them here—there will be duties held by individual and collective agents to bring about those reforms. Once the reforms are in place, the newly formed or reformed

collective agent will have duties based on the Dependency Principle to protect the population that is “people on Earth.”

The Coordination Principle also goes some way toward solving the problem of a lack of “political will” around R2P. Numerous R2P proponents (and detractors) cite a lack of will as one of the main reasons why many populations remain inadequately protected (Ban 2009, 15; G. Evans 2008, 223; Hehir 2012, 120–35; ICRtoP 2009, 5-6). States regularly engage in “inhumanitarian non-intervention,” that is, they lack the will to fulfil R2P intervention duties (Chesterman 2003). Presumably at least part of the reason for this is the high cost involved in fulfilling these duties, *if* an agent is to fulfil them independently (i.e., without coordination). However, if an agent could not protect a population via intervention on their own while realising positive expected value, then they do not have a duty to do on their own. Rather, each agent has an obligation to take protective measures only if sufficient other agents have signalled conditional willingness. So agents will not have duties to take on the full cost of the group action themselves. If others fail to signal, then an agent might still have an individual dependency duty, but this is likely to realise lower expected value—and so be less weighty—than is the sum of the duties to be responsive. (Of course, if agents lack political will to act no matter how many others act with them or no matter how cheap it is, then the Coordination Principle will have no practical effect on the issue of political will—though it will still generate duties.)

Finally note that R2P-related dependency and coordination duties will not co-exist, if the agent is unable to take the measures demanded by both principles. For example, suppose there are lots of things that various states, NGOs, and IGOs could do to reduce the likelihood that Syrians will suffer from gross harm in the next year. These agents may or may not have dependency duties to take these measures. For recall that dependency duties arise only if the agent’s *most efficacious* measures would realise *no less expected value*, regarding themselves and Syrians, than any other agent’s most efficacious measure. But the Syrians’ important interests might be more reliably fulfilled if that agent acted responsively to other agents to bring about some *even better* state of affairs for Syrians than the one that would be produced if the agent took their own individual measures. In this case, the most efficacious measures would not be the solo measures, but rather than multi-agent measures. Additionally, working with other agents would be less costly for the agent themselves. Thus, given that others have signalled conditional willingness, responsive measure will be more efficacious than individual measures. Because the individual measures are not the most efficacious measures, there will be no dependency duty. Instead, there will be a coordination duty to take the (more efficacious) responsive steps.

(Note that the principles can discount future benefits, if the best theory of costs and benefits dictates that they should. The question of whether they should is a question for the theory of interests with which we combine these principles, and it is a question on which this thesis—and, I take it, the R2P canon—takes no definitive stand. Thus a coordinating action that will have many positive effects, but only a hundred years from now, may or may not trump a solo action that will have fewer positive effects, but have them sooner.)

## 8.6 Conclusion

An adequate unifying, precisifying, moral justification of R2P has to explain three things: first, why (morally speaking) the home state has the primary duty; second, how to determine which particular agent(s) bear duties when the home state fails, and why this method of determination is the right one; and third, what the limits of the duties are. Together, the Dependency and Coordination Principles do this.

First, regarding the domestic duty's primacy, the Dependency Principle gives us the result that—as a matter of empirical fact—home states tend to have dependency duties to protect their own populations, with other states and international agents incurring duties only if the home state fails. It does not assert the domestic duty's primacy as a necessary truth. But then, we should not interpret the R2P canon or commentary as asserting that, either. The R2P is designed to deal with a political problem in the real world, in which there are good reasons to act as if (non-failed) states are, in *almost* all cases, best-placed.

Second, regarding the distribution of international duties, the Dependency and Coordination Principles together determine which agents have which duties, and they do this in a much more extensionally adequate way than self-interest, voluntary assumption, contribution, association, or proximity. That is, these other principles fail to produce duties where there intuitively should be duties; or they produce duties where there intuitively should not be duties. The Dependency and Coordination Principles, on the other hand, give us the right answers.

And these principles get these cases right *for the right reasons*. Many academic books and articles on R2P begin by describing terrible mass atrocities that have happened in the past, such as those in Rwanda in 1994, Bosnia in 1995, Kosovo in 1999, and Western Europe in the 1940s (e.g., Bellamy 2009, 1; G. Evans 2008, 1–2; Pattison 2010, 1–2; Thakur 2006, 1). The striking feature of these cases is the victims' plight, and the sharp moral imperative to alleviate that plight in whatever way is most effective. It is our terribly human vulnerability to vast political processes and powerful political agents that most animates

R2P proponents. It makes sense, then, that this vulnerability—and most especially, agents' being best-placed (whether alone or together) to respond to that vulnerability—should be the basis on which R2P duties are assigned. (In accordance with this, one might maintain that the “real” explanation of R2P is human rights—in which case, my argument should be read as stating that the best way to distribute the positive *duties* that come correlative to these human rights is along the lines of the Dependency and Coordination Principles.)

Third and finally, regarding the duties' limits, my two principles have scope for limiting R2P duties as our foundational theories and considered intuitions see fit. Exactly what gets counted as a benefit or a cost of a particular protective action, and exactly how much weight these benefits and costs are given, is in no way dictated by my principles. Additionally, the principles allow that strong defeaters might undermine the all-things-considered duties for reasons that are external to the logic of the duties themselves.

## **Chapter Nine:** **Conclusion**

In Chapter One, I introduced a basic thought: if you are dependent on someone in the right way, then that person has a moral duty. I laid out four sets of questions that we could ask about this basic thought, and asserted that this thesis would answer those questions in a way that shed light on various moral phenomena. I would now like to return to those questions, to lay bare my answers and the light that they have shed.

The first two sets of questions were about the *theory* of dependence-based duties. They were concerned with specifying the antecedent of the basic thought—that is, they were concerned with the conditions under which dependence generates duties. The first set of questions regarded the antecedent as it applies to agents acting on their own, while the second set of questions regarded the antecedent as it applies to agents acting together. I addressed these two sets of questions in Chapters Two and Three, producing the Dependency and Coordination Principles respectively. The questions were as follows.

First, over which interests do these duties range? Can there be dependence-based duties to provide others with anything whatsoever, or just the most important things? I answered that the duties range over important interests, but exactly which interests count as “important” was left substantially open. Whether an interest is sufficiently important to generate a duty in a given context depends on: the value of that interest’s being fulfilled; whether the extent of the interest’s importance is proportionate to the likelihood that it would be fulfilled if the right measures were taken; and how the interest’s value weighs up against the other costs and benefits that would be realised for the agent and dependent if measures were taken or not taken. These are questions for theories of wellbeing, and for theories of value more broadly. This answer demonstrated the ecumenicism of my account.

Second, how should we understand the ability, or capacity, that is required in order to have one of these duties? I suggested we think about the *measures* that an agent might take to fulfil the interest, where we assume that agents have control over whether they take measures. We should look at the measures they can take that are most likely to fulfil the interest. And we should ask: is the likelihood of that measure fulfilling the interest (if the measure is taken) proportionate to the importance of the interest, where important interests are proportionate to lower likelihoods? In other words, is it *likely enough* that the interest will be fulfilled if the measures is taken? If so, then the agent is “sufficiently capable” to bear a dependency duty to take the measure. Thus my answers to these first two questions fed into one another.

Third, what kinds of costs constrain these duties, and to what extent? The relevant costs are the net expected costs to the agent and to the dependent, since what we are examining is the relationship between *them*. I did not take a stand on exactly what kinds of benefits and costs to the agent and dependent matter here. Again, this was to be ecumenical among foundational theories. Part of the aim was to show just how much work dependence-based duties can do for us while retaining neutrality on a wide range of moral questions. Additionally, high costs to third parties might serve as a defeater, undermining the existence of an all-things-considered duty. But such a constraint would come from the outside, not from the internal logic of the dependence relation.

Finally, how must the duty-bearer's capacity compare to that of other agents—what does it mean to be “best-placed”? Skirting over some details to do with aggregating across cases, I answered that the best-placed agent is the one whose most efficacious measure for fulfilling a given important interest: is *sufficiently likely* to fulfil that interest; would, if taken, realise *positive* expected value regarding agent and dependent; and would, if taken, realise *no less* expected value (regarding agent and dependent) than any other agent's most efficacious measure for fulfilling that interest. Another way of looking at it is this: we should look at all measures (of all agents) that are *sufficiently capable* of fulfilling the interest and that have *positive* expected value regarding the dependent and the measure's agent, and then we should hone in on that one measure that has the *highest* expected value, regarding the dependent and the measure's agent, of all those measures. Whichever agent that measure belongs to is “best-placed.” These various answers combined to produce the Dependency Principle. This stated that when you are best-placed to fulfil an interest, then you have a duty to take the measure that renders you best-placed.

The second set of questions concerned dependence-based duties in group contexts. My answers shed light on important issues regarding group agency and shared intentions. These questions asked about the possibility and nature of group agency, group capacities, and group duties: in what sense might a group of individuals be best-placed to provide assistance, if none of them alone can exercise the capacity to assist? Can we talk about “the group” exercising the capacity (and having a duty to do so), and, if so, what does that mean? The answers to these questions depended crucially on whether the group in question was an agent—that is, whether it had a decision-making procedure that used distinct inputs, and had a distinct method for processing inputs to form outputs, from any one of its members, where members had committed to abiding by the group's decisions. Groups who meet these criteria can reliably produce multilateralism, and so have and exercise capacities that are more than the sum of their members' individual capacities. Groups who do not meet the agency criteria, on the other hand, do not have (and so

cannot exercise) capacities to assist that are additional to their members' several capacities to act responsively to one another.

How do groups' duties distribute to individuals? Again, the answer depends on whether the group is an agent. If the group is not an agent, then it makes no sense to speak of distribution of duties from the group to its members. If it is an agent, then the group can bear dependency duties under the Dependency Principle. The distribution of these duties is complex. When the group has a duty to see to it that X, then each member has a duty to act (as necessary) within their role to employ the group decision-making procedure to distribute roles to members such that: if enough members acted within their roles with a view to seeing to it that X (including, if other values allow it, cajoling, coercing, and covering for others), then that would be sufficient for it being the case that X in a high proportion of likely futures. Once these X-sufficient roles are distributed, each member has a duty to act within their role with a view to seeing to it that X.

The final question regarding group dependence-based duties asked: what duties might there be in cases where no such group exists but where a number of individuals could create a best-placed duty-bearing group or could in some other way together fulfil an important interest? The answer was: coordination duties. These are duties to either (i) take responsive steps with a view to a non-actual state-of-affairs in which an important interest is fulfilled, or (ii) take responsive steps towards there being a collective that can produce a non-actual state-of-affairs in which an important interest is fulfilled. As for dependency duties, the existence of coordination duties is sensitive to the expected value of coordination for the agents and dependent.

With these two sets of theoretical answers in place, we were able to turn to the *practice* of individual and collective dependence-based duties. Rather than taking on the impossible task of specifying the practice of dependence-based duties in general, I focused in on two particular doctrines: the ethics of care doctrine, which focuses (mainly) on individual, interpersonal ethics; and the Responsibility to Protect (R2P) doctrine, which focuses on collective, international ethics. Both of these are concerned with ethics as it plays out "on the ground," making them particularly apt for demonstrating the real-world upshots of dependence-based duties.

About the practice of individual dependence-based duties, I first asked: should we really be consciously entertaining dependence-based duties as we go about everyday life, or are they somehow self-effacing? I suggested that it is natural to think of them as being self-effacing, and that this allows us to resolve a tension within care ethics. This tension was that, on the one hand, care ethicists believe that ethical theory should positively endorse deliberation involving sympathy and direct attendance to concrete particulars, while, on the

other hand, care ethicists do posit general principles about rightness. It also allows us to use the principles to generate duties to have caring attitudes, where these attitudes would be undermined if the agent positively entertained the *duty* to do so.

The second question regarding individual practice was: how many of the duties recognised by common sense can be understood as instances of dependence-based duties? In my discussion of care ethics, I did not give a general answer to this question. Instead, I focused on those common sense duties that arise out of personal relationships. I considered the emotional (and material) interests that personal relatives fulfil for one another, and argued that these are the best explanation of the moral value of personal relationships and the duties that arise from them. These duties can thus be seen as dependence-based. I used this thought to argue that the Dependency and Coordination Principles will (perhaps only occasionally) entail duties to take steps to form new personal relationships.

Third, I asked whether dependence-based duties might call for attitudes and emotions, as well as for actions. I argued that objections to this idea are ultimately unconvincing. I therefore concluded that attitudes and emotions can be called for by dependence-based principles. In particular, the attitudes and emotions that are characteristic of “caring about” can be called for by the principles. The principles produce imperatives both to “care for” (roughly, to perform actions under the (perhaps only tacit) intention of fulfilling someone’s perceived interests) and to “care about” (roughly, to have attitudes that respond positively to the prospect of someone’s having a decent life).

The final set of questions concerned the *practice* of *collectives*’ dependence-based duties. I asked: can important large-scale real-world groups, such as states and intergovernmental organisations, bear dependence-based duties? I argued that they meet the criteria for bearing collective agency, and so that they can bear duties. However, for states, these duties may not distribute to individuals in quite the way one thinks: when a given state has a duty, which agents have distributed duties will be importantly dependent on that state’s decision-making procedure. Only agents who (*inter alia*) have a role in the state that affords them positive influence under the state’s decision-making procedure can bear distributed duties. This helps us to understand the sense in which ordinary citizens of democratic states are implicated in their states’ agency (and bear distributed duties when it has duties), while ordinary citizens of non-democratic states are not.

I then asked two final questions: to what extent are dependence-based duties borne out in actual political practice? How might they be fulfilled in international politics? To give specific answers these questions, I turned to the international political doctrine of R2P. This doctrine asserts a range of duties, and I argued that these duties are best understood as

dependence-based. Other explanations of R2P duties—in terms of self-interest, human rights, voluntary assumption, contribution, association, or proximity—either give the wrong answers on cases, or do not do enough to actually assign the duties to particular agents. Given this interpretation of R2P, the R2P canon and commentary have given us some answers to these last two questions. The answers seemed to be that there is much rhetorical support for dependence-based duties in the form of R2P, but that it remains to be seen the extent to which these duties will actually be borne out and implemented in international politics. Nonetheless, by viewing R2P duties through a dependence-based lens, I was able to clarify the contents and distribution of R2P duties.

There remain many interesting and important questions about dependence-based duties. Some questions can only be answered by appeal to foundational theories, or to broader theories of value—for example, we might wonder which other values can defeat dependence-based duties. Other questions can only be answered by delving into yet more real-world problems—for example, we might wonder what dependence-based duties have to say about our obligations to future generations. The persistence of these questions indicates that this thesis—on “the scope of dependence-based duties”—has not explored the full range of these duties’ scope. That scope is potentially huge, and to fully explore it would take a lifetime. But I hope nonetheless to have mapped important parts of its scope, and (perhaps more importantly) to have developed a framework for exploring the rest.

## Works Referenced

Abbott, Kenneth W. and Duncan Snidal. 2009. "Strengthening International Regulation Through Transnational New Governance: Overcoming the Orchestration Deficit." *Vanderbilt Journal of Transnational Law* 42(2), 501–578.

Abbott, Kenneth, Philipp Genschel, Duncan Snidal, and Bernhard Zangl. 2010. "IOs as Orchestrators." Draft paper presented to the SGIR 7th Pan-European International Relations Conference, Stockholm 9–11 September. Available at <<http://stockholm.sgir.eu/uploads/SGIR%20Stockholm%20100901.pdf>>, last accessed 20 June 2013.

Alexander, Jeffrey C. 1987. "The Social Requisites for Altruism and Voluntarism: Some Notes on What Makes a Sector Independent." *Sociological Theory* 5(2), 165–171.

Anderson, Elizabeth. 1993. *Value in Ethics and Economics*. Cambridge, Mass.: Harvard University Press.

Annan, Kofi. 1999a. "Two Concepts of Sovereignty." *The Economist* September 16. Available at <<http://www.economist.com/node/324795>>, last accessed 20 June 2013.

Annan, Kofi. 1999b. "Address Presenting Annual Report to the Opening Meeting of the United Nations General Assembly." United Nation Press Release SG/SM/7136 GA/9596, 20 September. Available at <<http://www.un.org/News/Press/docs/1999/19990920.sgsm7136.html>>, last accessed 20 June 2013.

Annan, Kofi. 2000. *We the Peoples: The Role of the United Nations in the 21st Century*. New York: United Nations.

Arbour, Louise. 2008. "The Responsibility to Protect as a Duty of Care in International Law and Practice." *Review of International Studies* 34(3), 445–458.

Arrow, Kenneth. 1963. *Social Choice and Individual Values*. Second edition. New Haven: Yale University Press.

Ashford, Elizabeth. 2003. "The Demandingness of Scanlon's Contractualism." *Ethics* 113(2), 273–302.

Bagnoli, Carla. 2004. "Humanitarian Intervention as a Perfect Duty: A Kantian Argument." *Nomos* 47, 117–148.

Baier, Annette C. 1981. "Cartesian Persons." *Philosophia* 10(3–4), 169–188.

Ban, Ki-Moon. 2009. *2009 Report: Implementing the Responsibility to Protect*. United Nations General Assembly Report of the Secretary-General. January 12. Available at <[http://responsibilitytoprotect.org/SGRtoPEng%20\(4\).pdf](http://responsibilitytoprotect.org/SGRtoPEng%20(4).pdf)>, last accessed 20 June 2013.

Baron, Marcia. 1984. "The Alleged Moral Repugnance of Acting from Duty." *Journal of Philosophy* 81(4), 197–220.

Baron, Marcia. 1991. "Impartiality and Friendship." *Ethics* 101(4), 836–857.

Barry, Christian and Gerhard Øverland. 2013. "How Much for the Child?" *Ethical Theory and Moral Practice* 16(1), 189–204.

Barry, Christian and Nicholas Southwood. 2011. "What Is Special About Human Rights?" *Ethics & International Affairs* 25(3), 369–383.

Beitz, Charles. 1979. *Political Theory and International Relations*. Princeton: Princeton University Press.

Beitz, Charles. 2009. *The Idea of Human Rights*. Oxford: Oxford University Press.

Bellamy, Alexander. 2008. "The Responsibility to Protect." Pp. 422–437 in P.D. Williams (ed.), *Security Studies: An Introduction*. London; New York: Routledge.

Bellamy, Alex J. 2009. *The Responsibility to Protect: The Global Effort to End Mass Atrocities*. Cambridge; Malden: Polity.

Bellamy, Alex J. 2010. "The Responsibility to Protect: Five Years On." *Ethics & International Affairs* 24(2), 143–169.

Benhabib, Seyla. 1987. "The Generalized and the Concrete Other." Pp. 154–177 in E.F. Kittay and D. Meyers (eds), *Women and Moral Theory*. Totowa: Rowman & Littlefield.

Bhandary, Asha. 2010. "Dependency in Justice: Can Rawlsian Liberalism Accommodate Kittay's Dependency Critique?" *Hypatia* 25(1), 140–156.

Blum, Lawrence. 1980. *Friendship, Altruism, and Morality*. London: Routledge and Kegan Paul.

Bratman, Michael. 1992. "Shared Cooperative Activity." *Philosophical Review* 101(2), 327–341.

Brennan, Geoffrey and Nicholas Southwood. 2007. "Feasibility in Action and Attitude." Published online only in T. Ronnow-Rasmussen, B. Petersson, J. Jonefsson, and D. Egonsson (eds), *Hommage a Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. Available at <[www.fil.lu.se/hommageawlodek](http://www.fil.lu.se/hommageawlodek)>, last accessed 20 June 2013.

Brock, Gillian and Harry Brighouse (eds). 2005. *The Political Philosophy of Cosmopolitanism*. Cambridge: Cambridge University Press.

Brown, Chris. 2001. "Moral Agency and International Society." *Ethics & International Affairs* 15(2), 87–98.

Bubeck, Diemut. 1995. *Care, Gender and Justice*. Oxford: Clarendon Press.

Buchanan, Allen. 1999. "The Internal Legitimacy of Humanitarian Intervention. *Journal of Political Philosophy* 7(1), 71–87.

Bureau of Ocean Energy Management, Regulation and Enforcement (BOEMRE). 2011. *Report Regarding the Causes of the April 20, 2010 Macondo Well Blowout*. Available online at <[http://docs.lib.noaa.gov/noaa\\_documents/DWH\\_IR/reports/dwhfinal.pdf](http://docs.lib.noaa.gov/noaa_documents/DWH_IR/reports/dwhfinal.pdf)>, last accessed 20 June 2013.

Butt, Daniel. 2007. "On Benefiting from Injustice." *Canadian Journal of Philosophy* 37(1), 129–152.

Calhoun, Cheshire. 1988. "Justice, Care, Gender Bias." *Journal of Philosophy* 85(9), 451–463.

Caney, Simon. 2005. *Justice Beyond Borders: A Global Political Theory*. Oxford: Oxford University Press.

Carment, David and Martin Fischer. 2009. "R2P and the Role of Regional Organisations in Ethnic Conflict Management, Prevention and Resolution: The Unfinished Agenda." *Global Responsibility to Protect* 1, 261–290.

Chappell, Timothy (ed.). 2009. *The Problem of Moral Demandingness: New Philosophical Essays*. London: Palgrave Macmillan.

Chatterjee, Deen (ed.). 2004. *The Ethics of Assistance*. Cambridge: Cambridge University Press.

Chesterman, Simon. 2003. "Hard Cases Make Bad Law." Pp. 46–61 in A. Lang (ed.), *Just Intervention*. Washington, DC: Georgetown University Press.

Chicago Council on Global Affairs and WorldPublicOpinion.org. 2007. "World Publics Favor New Powers for the UN." May 9. Available at <<http://www.worldpublicopinion.org/pipa/articles/btunitednationsra/355.php>>, last accessed June 20 2013.

Chisholm, Roderick. 1963. "Supererogation and Offense: A Conceptual Scheme for Ethics." *Ratio* 5(1), 1–14.

Chuard, Philippe and Nicholas Southwood. 2009. "Epistemic Norms Without Voluntary Control." *Noûs* 43(4), 599–632.

Collins, Stephanie and Holly Lawford-Smith. 2013. "The Transfer of Duties: From Individuals to States and Back Again." Unpublished manuscript.

Collins, Stephanie. 2013. "Collectives' Duties and Collectivization Duties." *Australasian Journal of Philosophy* 91(2), 231–248.

Collins, Stephanie. Forthcoming. "Duties to Make Friends." *Ethical Theory and Moral Practice*, BSET Special Issue.

Cottingham, John. 2010. "Impartiality and Ethical Formation." Pp. 65–83 in B. Feltham and J. Cottingham (eds), *Partiality and Impartiality: Morality, Special Relationships and the Wider World*. Oxford: Oxford University Press.

Crawford, James. 2012. *Brownlie's Principles of Public International Law*. Eighth edition. Oxford: Oxford University Press.

Cullity, Garrett. 2000. "Pooled Beneficence." Pp. 9–42 in M.J. Almeida (ed.), *Imperceptible Harms and Benefits*. Dordrecht: Kluwer.

Cullity, Garrett. 2004a. *The Moral Demands of Affluence*. New York: Oxford University Press.

Cullity, Garrett. 2004b. "Sympathy, Discernment and Reasons." *Philosophy and Phenomenological Research* 68(1), 37–62.

Cullity, Garrett. 2009. "Demandingness and Arguments from Presupposition." Pp. 8–34 in T. Chappell (ed.), *The Problem of Moral Demandingness: New Philosophical Essays*. London: Palgrave Macmillan.

Dancy, Jonathan. 2004. *Ethics without Principles*. Oxford: Oxford University Press.

Darwall, Stephen. 2002. *Welfare and Rational Care*. Princeton: Princeton University Press.

Deng, Francis, Sadikiel Kimaro, Terrence Lyons, Donald Rothchild, and I. William Zartman. 1996. *Sovereignty as Responsibility: Conflict Management in Africa*. Washington: The Brookings Institution.

DeSousa, Ronald. "Emotion." In E. Zalta (ed.), *Stanford Encyclopaedia of Philosophy*. Available at <<http://plato.stanford.edu/archives/spr2012/entries/emotion/>>, last accessed 20 June 2013.

Driver, Julia. 2005. "Consequentialism and Feminist Ethics." *Hypatia* 20(4), 183–199.

Elster, Jon. 1985. *Sour Grapes: Studies in the Subversion of Rationality*. Cambridge: Cambridge University Press.

Engster, Daniel. 2005. "Rethinking Care: The Practice of Caring and the Obligation to Care." *Hypatia* 23(3), 50–74.

Engster, Daniel. 2007. *The Heart of Justice: A Political Theory of Caring*. New York: Oxford University Press.

Erskine, Toni. 2001. "Assigning Responsibilities to Institutional Moral Agents: the Case of States and Quasi-states." *Ethics & International Affairs* 15(2), 67–89.

Estlund, David. 2011. "Human Nature and the Limits (If Any) of Political Philosophy." *Philosophy & Public Affairs* 39(3), 207–237.

Evans, Dylan. 2002. "The Search Hypothesis of Emotions." *British Journal of the Philosophy of Science* 53, 497–509.

Evans, Gareth. 2008. *The Responsibility to Protect: Ending Mass Atrocity Crimes Once and For All*. Washington: The Brookings Institution.

Feinberg, Joel. 1968. "Collective Responsibility." *Journal of Philosophy* 65, 674–688.

Feldman, Fred. 1999. "Comments on *Living High and Letting Die*." *Philosophy and Phenomenological Research* 59(1), 195–201.

Feltham, Brian and John Cottingham (eds). 2010. *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*. Oxford: Oxford University Press.

Fine, Kit. 2012. "Response Dependence and Joint Intention." Paper presented at Collective Intentionality VIII, University of Manchester, 31 August 2012.

Fisher, Berenice and Joan Tronto. 1990. "Toward a Feminist Theory of Care." Pp. 35–62 in E. Abel and M. Nelson (eds), *Circles of Care: Work and Identity in Women's Lives*. Albany: State University of New York Press.

Frankfurt, Harry G. 1982. "Comments on MacIntyre." *Synthese* 53(2), 319–321.

Frankfurt, Harry G. 1999. *Necessity, Volition and Love*. Cambridge: Cambridge University Press.

French, Peter A. 1979. "The Corporation as a Moral Person." *American Philosophical Quarterly* 16(3), 207–215.

Friedman, Marilyn. 1998. "Impartiality." Pp. 393–401 in A.M. Jagger and I.M. Young (eds), *A Companion to Feminist Philosophy*. Oxford: Blackwell.

Friedman, Marilyn. 1991. "The Practice of Partiality." *Ethics* 101(4), 818–835.

Gilbert, Margaret. 2001. "Collective Preferences, Obligations, and Rational Choice." *Economics and Philosophy* 17(1), 109–119.

Gilbert, Margaret. 2006. "Who's to Blame? Collective Moral Responsibility and its Implications for Group Members." *Midwest Studies in Philosophy* 30(1), 94–114.

Gilbert, Pablo and Holly Lawford-Smith. 2012. "Political Feasibility: A Conceptual Exploration." *Political Studies* 60(4), 809–825.

Gilligan, Carol. 1982. *In A Different Voice*. Cambridge, MA: Harvard University Press.

Global Centre for the Responsibility to Protect (GCR2P). 2009. "Implementing the Responsibility to Protect: The 2009 General Assembly Debate: An Assessment." Ralph Bunche Institute for International Studies, The CUNY Graduate Center. Available at <<http://www.globalr2p.org/publications/42>>, last accessed 20 June 2013.

Goldman, Holly. 1978. "Doing the Best One Can." Pp. 185–214 in A. Goldman and J. Kim (eds), *Values and Morals*. Dordrecht: Reidel.

Goodin, Robert E. 1985. *Protecting the Vulnerable*. Chicago: University of Chicago Press.

Goodin, Robert E. 1988. "What is so Special About Our Fellow Countrymen?" *Ethics* 98(4), 663–686.

Goodin, Robert E. 2003. "Democratic Accountability: the Distinctiveness of the Third Sector." *European Journal of Sociology* 44, 359–396.

Goodin, Robert E. 1995. *Utilitarianism as a Public Philosophy*. Cambridge: Cambridge University Press.

Goodin, Robert E. 2009. "Demandingness as a Virtue." *Journal of Ethics* 13, 1–13.

Goodin, Robert E. 2012. "Excused by the Unwillingness of Others?" *Analysis* 72(1), 18–24.

Grant, Thomas D. 1999. *The Recognition of States: Law and Practice in Debate and Evolution*. Westport: Praeger.

Griffiths, Paul. 1997. *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.

Grigorescu, Alexandru. 2007. "Transparency of Intergovernmental Organizations: The Roles of Member-States, International Bureaucracies and Non-Governmental Organizations." *International Studies Quarterly* 51(3), 625–648.

Harbour, Frances. 2004. "Moral Agency and Moral Responsibility in Humanitarian Intervention." *Global Society* 18(1), 61–75.

Hehir, Aidan. 2012. *The Responsibility to Protect: Rhetoric, Reality and the Future of Humanitarian Intervention*. Basingstoke: Palgrave Macmillan.

Held, Virginia. 1970. "Can a Random Collective Be Morally Responsible?" *Journal of Philosophy* 67, 471–481.

Held, Virginia. 1987. "Feminism and Moral Theory." Pp. 111–128 in E. F. Kittay and D. Meyers (eds), *Women and Moral Theory*. Totowa: Rowman & Littlefield.

Held, Virginia. 1990. "Feminist Transformations of Moral Theory." *Philosophy and Phenomenological Research* 50, 321–344.

Held, Virginia. 1993. *Feminist Morality*. Chicago: University of Chicago Press.

Held, Virginia (ed.). 1995. *Justice and Care: Essential Readings in Feminist Ethics*. Boulder: Westview.

Held, Virginia. 2004. "Taking Care: Care as Practice and Value." Pp. 59–71 in C. Calhoun (ed.), *Setting the Moral Compass: Essays by Women Philosophers*. Oxford: Oxford University Press.

Held, Virginia. 2006. *The Ethics of Care: Personal, Political, and Global*. Oxford: Oxford University Press.

Held, Virginia. 2008. "Military Intervention and the Ethics of Care." *Southern Journal of Philosophy* 46(S1), 1–20.

High-Level Panel on Threats, Challenges, and Change (HLPTCC). 2004. *A More Secure World: Our Shared Responsibility*. United Nations Department of Public Information.

Hooker, Brad and Margaret Little (eds). 2000. *Moral Particularism*. Oxford: Oxford University Press.

Hooker, Brad. 1999. "Sacrificing for the Good of Strangers—Repeatedly." *Philosophy and Phenomenological Research* 59(1), 177–181.

Hume, David. 1987[1742]. *Essays, Moral, Political, And Literary*. E.F. Miller (ed.). Indianapolis, IN: Liberty Fund, Inc.

Hursthouse, Rosalind. 1999. *On Virtue Ethics*. Oxford: Oxford University Press.

Igneski, Violetta. 2006. "Perfect and Imperfect Duties of Aid." *Social Theory and Practice* 32, 339–366.

International Coalition for the Responsibility to Protect (ICRtoP). 2009. "Report on the General Assembly Plenary Debate on the Responsibility to Protect." September 15. Available at <[http://www.responsibilitytoprotect.org/ICRtoP%20Report-General\\_Assembly\\_Debate\\_on\\_the\\_Responsibility\\_to\\_Protect%20FINAL%209\\_22\\_09.pdf](http://www.responsibilitytoprotect.org/ICRtoP%20Report-General_Assembly_Debate_on_the_Responsibility_to_Protect%20FINAL%209_22_09.pdf)>, accessed 20 June 2013.

International Commission on Intervention and State Sovereignty (ICISS). 2001. *The Responsibility to Protect: Report of the International Commission on Intervention and State Sovereignty*. Ottawa: International Development Research Centre.

International Law Commission (ILC). 2011. *Draft Articles on the Responsibility of International Organizations*. Available at <[http://untreaty.un.org/ilc/texts/instruments/english/draft%20articles/9\\_11\\_2011.pdf](http://untreaty.un.org/ilc/texts/instruments/english/draft%20articles/9_11_2011.pdf)>, last accessed 24 June 2013.

International Monetary Fund (IMF). 2013. *IMF Quotas*. Available at <<http://www.imf.org/external/np/exr/facts/quotas.htm>>, last accessed 24 June 2013.

Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101(3), 461–482.

Jaggar, Alison. 1995. "Caring as a Feminist Practice of Moral Reason." Pp. 179–202 in V. Held (ed.), *Justice and Care: Essential Readings*, Boulder: Westview Press.

James, Scott. 2007. "Good Samaritans, Good Humanitarians." *Journal of Applied Philosophy* 24(3), 238–254.

Jasentuliyana, Nandasiri (ed.). 1995. *Perspectives on International Law*. London: Kluwer Law International.

Jensen, Mark. 2009. "The Limits of Practical Possibility." *Journal of Political Philosophy* 17(2), 168–184.

Kamm, Frances M. 2000. "Does Distance Matter Morally to the Duty to Rescue?" *Law and Philosophy* 19(6), 655–681.

Keller, Simon. 2006. "Four Theories of Filial Duty." *Philosophical Quarterly* 56(223), 254–274.

Kittay, Eva Feder. 1996. "Human Dependency and Rawlsian Equality." Pp. 219–266 in D.T. Meyers (ed.), *Feminist Rethink the Self*. Colorado: Westview Press.

Kittay, Eva Feder. 1999. *Love's Labor: Essays on Women, Equality and Dependence*. New York: Routledge.

Kittay, Eva Feder. 2005. "Dependency, Difference and the Global Ethic of Longterm Care." *Journal of Political Philosophy* 13(4), 443–469.

Kittay, Eva Feder. 2008. "The Global Heart Transplant and Caring across National Boundaries." *Southern Journal of Philosophy* 46(S1), 138–165.

Koggel, Christine and Joan Orme. 2010. "Care Ethics: New Theories and Applications." *Ethics and Social Welfare* 4(2), 109–114.

Kolodny, Niko. 2003. "Love as Valuing a Relationship." *Philosophical Review* 112(2), 135–189.

Kolodny, Niko. 2010a. "Which Relationships Justify Partiality? General Considerations and Problem Cases." Pp. 169–193 in B. Feltham and J. Cottingham (eds), *Partiality and Impartiality: Morality, Special Relationships and the Wider World*. Oxford: Oxford University Press.

Kolodny, Niko. 2010b. "Which Relationships Justify Partiality? The Case of Parents and Children." *Philosophy & Public Affairs* 38, 37–75.

Korsgaard, Christine. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.

Kumar, Rahul. 1999. "Defending the Moral Moderate: Contractualism and Common Sense." *Philosophy & Public Affairs* 28, 275–309.

Kutz, Christopher. 2000. *Complicity*. Cambridge: Cambridge University Press.

Lawford-Smith, Holly. 2012. "The Feasibility of Collectives' Actions." *Australasian Journal of Philosophy* 90(3), 453–467.

Liao, Matthew S. 2006. "The Right of Children to be Loved." *Journal of Political Philosophy* 14(4), 420–440.

Lichtenberg, Judith. 2004. "Absence and the Unfond Heart: Why People Are Giving Less Than They Might Be." Pp. 75–100 in D. Chatterjee (ed.), *The Ethics of Assistance: Morality and the Distant Needy*. New York: Cambridge University Press.

List, Christian and Philip Pettit. 2011. *Group Agency*. Oxford: Oxford University Press.

Luck, Edward. 2002. "Prevention: Theory and Practice." Pp. 251–271 in F. Osler Hampson and D. Malone (eds), *From Reaction to Conflict Prevention: Opportunities for the UN System*. Boulder: Lynne Rienner.

Mackenzie, Catriona and Natalie Stoljar (eds). 2000. *Relational Autonomy: Feminist Perspective on Autonomy, Agency and the Social Self*. New York: Oxford University Press.

May, Larry and Robert Strikwerda. 1994. "Men in Groups: Collective Responsibility for Rape." *Hypatia* 9, 134–151.

May, Larry. 1992. *Sharing Responsibility*. Chicago: University of Chicago Press.

McMahan, Jeff. 1996. "Intervention and Collective Self-Determination." *Ethics & International Affairs* 10, 1–24.

Menon, Rajan. 2009. "Pious Words Puny Deeds: The 'International Community' and Mass Atrocities." *Ethics & International Affairs* 23(3), 235–245.

Meyers, Diana T. 1987. "The Socialized Individual and Individual Autonomy: An Intersection between Philosophy and Psychology." Pp. 139–153 in E. F. Kittay and D. Meyers (eds), *Women and Moral Theory*. Totowa: Rowman & Littlefield.

Miller, David. 2001. "Distributing Responsibilities." *Journal of Political Philosophy* 9(4), 453–471.

Miller, David. 2004. "Holding Nations Responsible." *Ethics* 114, 240–268.

Miller, David. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.

Miller, Richard. 2004. "Beneficence, Duty, and Distance." *Philosophy & Public Affairs* 32, 357–383.

Miller, Sarah Clark. 2005. "Need, Care and Obligation." *Royal Institute of Philosophy* 80, supp. 57 (supplement on "The Philosophy of Need"), 137–160.

Miller, Sarah Clark. 2010. "Cosmopolitan Care." *Ethics and Social Welfare* special issue on "Care Ethics: New Theories and Applications" 4(2), 145–157.

Murphy, Liam B. 2000. *Moral Demands in Nonideal Theory*. New York: Oxford University Press.

Murphy, Liam. 1993. "The Demands of Beneficence." *Philosophy & Public Affairs* 22(4), 267–292.

Narveson, Jan. 2003. "We Don't Owe Them a Thing! A Tough-Minded but Soft-Hearted View of Aid to the Faraway Needy." *The Monist* 86, 419–433.

Nedelsky, Jennifer. 1989. "Reconceiving Autonomy: Sources, Thoughts and Possibilities." *Yale Journal of Law and Feminism* 1, 7–36.

Nickel, James. 2010. "Human Rights." In E. Zalta (ed.), *Stanford Encyclopaedia of Philosophy*.

Available at <<http://plato.stanford.edu/entries/human-rights>>, last accessed 20 June 2013.

Noddings, Nel. 1984. *Caring: A Feminine Approach to Ethics and Moral Education*. Berkeley: University of California Press.

Noddings, Nel. 1995. "Caring." Pp. 7-30 in V. Held (ed.), *Justice and Care: Essential Readings in Feminist Ethics*. Boulder: Westview Press.

Noddings, Nel. 2002. *Starting at Home: Caring and Social Policy*. Los Angeles: University of California Press.

Noddings, Nel. 1999. "Two Concepts of Caring." *Philosophy of Education* 1999, 36–39.

Nussbaum, Martha C. 2006. *Frontiers of Justice: Disability, Nationality, Species Membership*. Cambridge, MA: Belknap Press.

Nussbaum, Martha. 2003. *Upheavals of Thought: The Intelligence of Emotions*. Cambridge: Cambridge University Press.

O'Neill, Onora. 2005. "The Dark Side of Human Rights." *International Affairs* 81, 427–439.

Oakley, Justin. 1992. *Morality and the Emotions*. London: Routledge.

Office of the President of the General Assembly (OPGA). 2009. "Concept Note on Responsibility to Protect Populations from Genocide, War Crimes, Ethnic Cleansing and Crimes Against Humanity." New York: Office of the President of the General Assembly, United Nations. Available at <<http://www.un.org/ga/president/63/interactive/protect/conceptnote.pdf>>, last accessed 20 June 2013.

Otsuka, Michael. 1991. "The Paradox of Group Beneficence." *Philosophy & Public Affairs* 20, 132–149.

Oxford English Dictionary Online. Oxford University Press. Available at <<http://www.oed.com>>, last accessed 20 June 2013.

Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.

Pasternak, Avia. Forthcoming. "Limiting States' Corporate Responsibility." *Journal of Political Philosophy*.

Pattison, James. 2010. *Humanitarian Intervention and the Responsibility to Protect: Who Should Intervene?* Oxford: Oxford University Press.

Pettit, Philip and David Schweikard. 2006. "Joint Actions and Group Agents." *Philosophy of the Social Sciences* 36, 18–39.

Pettit, Philip and Robert Goodin. 1986. "The Possibility of Special Duties." *Canadian Journal of Philosophy* 16(4), 651–676.

Pettit, Philip. 1997. "Love and its Place in Moral Discourse." Pp. 153–163 in R. Lamb (ed.), *Love Analyzed*. Boulder: Westview Press.

Pettit, Philip. 2003. "Akrasia, Collective and Individual." Pp. 68–96 in S. Stroud and C. Tappolet (eds), *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press.

Pettit, Philip. 2007. "Responsibility Incorporated." *Ethics* 117, 141–201.

Pettit, Philip. 2010. "Groups with Minds of Their Own." Pp. 167–193 in A. Goldman and D. Whitcomb (eds), *Social Epistemology: Essential Readings*. Oxford: Oxford University Press.

Pogge, Thomas. 2002. *World Poverty and Human Rights*. Cambridge: Polity Press.

Prinz, Jesse. Forthcoming. "Is Empathy Necessary for Morality?" In P. Goldie and A. Coplan (eds), *Empathy: Philosophical and Psychological Perspectives*. Oxford: Oxford University Press. Available at <<http://subcortex.com/>>, last accessed 20 June 2013.

Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy & Public Affairs* 13(2), 134–171.

Rainbolt, George. 2000. "Perfect and Imperfect Obligations." *Philosophical Studies* 98(3), 233–256.

Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Oxford University Press.

Robinson, Fiona. 1999. *Globalizing Care: Ethics, Feminist Theory, and International Relations*. Boulder: Westview Press.

Robinson, Fiona. 2010. "After Liberalism in World Politics? Towards an International Political Theory of Care." *Ethics and Social Welfare* 4(2), 130–144.

Robinson, Fiona. 2011. *The Ethics of Care: A Feminist Approach to Human Security*. Philadelphia: Temple University Press.

Rodin, David. 2006. "The Responsibility to Protect and the Logic of Rights." Pp. 45–60 in O. Jütersonke and K. Krause (eds), *From Rights to Responsibilities*. Geneva: Programme for Strategic and International Security Studies.

Romain, Dianne. 1992. "Care and Confusion." Pp. 27–44 in E.B. Cole and S. Coultrap-McQuin (eds), *Explorations in Feminist Ethics*. Bloomington: Indiana University Press.

Ruddick, Sara. 1980. "Maternal Thinking." *Feminist Studies* 6(2), 342–367.

Ruddick, Sara. 1989. *Maternal Thinking: Toward a Politics of Peace*. Boston: Beacon Press.

Ruddick, Sara. 1992. "From Maternal Thinking to Peace Politics." Pp. 141–155 in E. Browning-Cole and S. Coultrap-McQuin (eds), *Explorations in Feminist Ethics*. Bloomington: Indiana University Press.

Ruddick, Sara. 1998. "Care as Labor and Relationship." Pp. 3-25 in J.C. Haber and M.S. Halfon (eds), *Norms and Values: Essays on the Work of Virginia Held*. Lanham: Rowman and Littlefield.

Ruddick, Sara. 2002. "An Appreciation of *Love's Labour*." *Hypatia* 17(3), 214–224.

Runciman, David. 2012. "What is Realistic Political Philosophy?" *Metaphilosophy* 43(1–2), 58–70.

Sandel, Michael. 1982. *Liberalism and the Limits of Justice*. Cambridge: Cambridge University Press.

Scanlon, T.M. 1975. "Preference and Urgency." *Journal of Philosophy* 72(19), 655–669.

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

Scheffler, Samuel. 1997. "Relationships and Responsibilities." *Philosophy & Public Affairs* 26, 189–209.

Scheffler, Samuel. 2001. *Boundaries and Allegiances: Problems of Justice and Responsibility in Liberal Thought*. Oxford: Oxford University Press.

Scheffler, Samuel. 2010. *Equality and Tradition: Questions of Value in Moral and Political Theory*. Oxford: Oxford University Press.

Shaw, Malcolm Nathan. 2003. *International Law*. Cambridge: Cambridge University Press.

Shue, Henry. 1996. *Basic Rights: Subsistence, Affluence and US Foreign Policy*. Second edition. Princeton: Princeton University Press.

Sidgwick, Henry. 1907. *Methods of Ethics*. Seventh edition. Indianapolis: Hackett.

Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1(3), 229–243.

Singer, Peter. 2009. *The Life You Can Save*. New York: Random House.

Slote, Michael. 1998. "The Justice of Caring." *Social Philosophy and Policy* 15(1), 171–195.

Slote, Michael. 1999. "Caring Versus the Philosophers." *Philosophy of Education* 1999, 25–35.

Smith, Holly. 2010. "Subjective Rightness." *Social Philosophy & Policy* 27(2), 64–110.

Sobel, David. 2007. "The Impotence of the Demandingness Objection." *Philosophers' Imprint* 7(8), 1–17.

Solomon, Robert C. (ed.). 2004. *Thinking about Feeling: Contemporary Philosophers on Emotions*. New York: Oxford University Press.

Sonderholm, Jorn. Forthcoming. "World Poverty, Positive Duties and the Over-demandingness Objection." *Politics, Philosophy & Economics*.

Steffek, Jens. 2010. "Public Accountability and the Public Sphere of International Governance." *Ethics & International Affairs* 24(1), 45–68.

Stocker, Michael. 1971. "Ought and Can." *Australasian Journal of Philosophy* 49(3), 303–316.

Stocker, Michael. 1976. "The Schizophrenia of Modern Ethical Theories." *The Journal of Philosophy* 73(14), 453–466.

Stocker, Michael. 1996. *Valuing Emotions*. Cambridge: Cambridge University Press.

Strawson, Peter. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48, 1–25.

Stueber, Karsten. 2008. "Empathy." In E. Zalta (ed.), *Stanford Encyclopaedia of Philosophy*. Available at <<http://plato.stanford.edu/archives/fall2008/entries/empathy/>>, last accessed 20 June 2013.

Sudgen, Robert. 2000. "Team Preferences." *Economics & Philosophy* 16(2), 175–204.

Tan, Kok-Chor. 2005. "The Duty to Protect." Pp. 84–116 in T. Nardin and M. Williams (eds), *Humanitarian Intervention. NOMOS ILVII*. New York: New York University Press.

Thakur, Ramesh. 2006. "The United Nations, Peace and Security: From Collective Security to the Responsibility to Protect." Cambridge: Cambridge University Press.

Tronto, Joan. 1987. "Beyond Gender Difference to a Theory of Care." *Signs* 12(4), 644–663.

Tronto, Joan. 1993. *Moral Boundaries: A Political Argument for an Ethic of Care*. London: Routledge.

Tronto, Joan. 1995. "Beyond Caring: What Can Feminists Learn About Morality from Caring?" Pp. 101-116 in V. Held (ed.), *Justice and Care: Essential Readings in Feminist Ethics*. Boulder: Westview Press.

Tronto, Joan. 2010. "Creating Caring Institutions: Politics, Plurality, and Purpose." *Ethics and Social Welfare* 4(2), 158–171.

Tuomela, Raimo. 2006. "Joint Intention, We-Mode and I-Mode." *Midwest Studies in Philosophy* XXX, 35–58.

Unger, Peter. 1996. *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.

Unger, Peter. 1999. "Replies." *Philosophy and Phenomenological Research* 59(1), 203–216.

United Nations Department of Public Information (UNDPI). 2012a. "World not Fulfilling 'Never Again' Vow, Secretary-General Tells General Assembly Meeting in Responsibility to Protect." UNDPI News and Media Division, New York. 5 September. Available at <<http://www.un.org/News/Press/docs/2012/ga11270.doc.htm>>, last accessed 20 June 2013.

United Nations Department of Public Information (UNDPI). 2012b. "Secretary-General, Addressing Drake University Students on 'Global Citizenship', says no Issue can

be Seen in Isolation as World Changes.” UNDPI News and Media Division, New York. 5 September. Available at <<http://www.un.org/News/Press/docs/2012/sgsm14595.doc.htm>>, last accessed 25 June 2013.s

United Nations General Assembly (UNGA). 2005. *Resolution Adopted by the General Assembly: 2005 World Summit Outcome*. October 24. Available at <<http://unpan1.un.org/intradoc/groups/public/documents/un/unpan021752.pdf>>, last accessed 20 June 2013.

United Nations Security Council (UNSC). 2006. “Protection of Civilians in Armed Conflict.” Resolution 1674. Available at <<http://daccess-dds-ny.un.org/doc/UNDOC/GEN/N06/331/99/PDF/N0633199.pdf?OpenElement>>, last accessed 20 June 2013.

United Nations. 2009. *Fact Sheet: United Nations Peacekeeping*. Available at <<http://www.un.org/en/peacekeeping/documents/factsheet.pdf>>, last accessed June 24 2013.

Vranas, Peter B.M. 2007. “I Ought, Therefore I Can.” *Philosophical Studies* 136(2), 167–216.

Walker, Margaret Urban. 1989. “Moral Understandings: Alternative ‘Epistemology’ for a Feminist Ethics.” *Hypatia* 4, 15–28.

Walzer, Michael. 1983. *Spheres of Justice: A Defense of Pluralism and Equality*. New York: Basic Books.

Walzer, Michael. 2006. *Just and Unjust Wars*. Fourth edition. New York: Basic Books.

Weber, Max. 1919[1946]. *Politics as a Vocation*. Pp. 77–128 in H.H. Gerth and C.W. Mills (eds and trans), *From Max Weber: Essays in Sociology*. New York: Oxford University Press.

Wedgwood, Ruth. 2009. “Democracies, Human Rights, and Collective Action.” *Ethics & International Affairs* 23(1), 27–37.

Wellman, Christopher. 1997. “Associative Allegiances and Political Obligations.” *Social Theory and Practice* 23, 181–204.

Welsh, Jennifer and Maria Banda. 2010. “International Law and the Responsibility to Protect: Clarifying or Expanding States’ Responsibilities?” *Global Responsibility to Protect* 2, 213–231.

Welsh, Jennifer. 2007. “The Responsibility to Protect: Securing the Individual in International Society.” Pp. 363–384 in B. Goold and L. Lazarus (eds), *Security and Human Rights*. Oxford: Hart.

Welsh, Jennifer. 2009. “Implementing the Responsibility to Protect.” Policy Brief 1/2009, Oxford Institute for Ethics, Law and Armed Conflict. Available at 258

<[http://www.elac.ox.ac.uk/downloads/r2p\\_policybrief\\_180209.pdf](http://www.elac.ox.ac.uk/downloads/r2p_policybrief_180209.pdf)>, last accessed 20 June 2013.

Wendt, Alexander. 2004. "The State as Person in International Theory." *Review of International Studies* 30(2), 289–316.

Wolf, Susan. 1982. "Moral Saints." *Journal of Philosophy* 79, 419–439.

Woods, Ngaire. 2003. "Holding Intergovernmental Institutions to Account." *Ethics & International Affairs* 17(1), 69–80.

Wringe, Bill. 2010. "Global Obligations and the Agency Objection." *Ratio* 23(2), 217–231.

Young, Iris Marion. 1990. *Justice and the Politics of Difference*. Princeton: Princeton University Press.