

ABSTRACT

Title of Dissertation: REASONS FOR RESCUE:
AN ESSAY ON BENEFICENCE

Scott M. James, Doctor of Philosophy, 2005

Dissertation directed by: Professor Christopher W. Morris
Department of Philosophy

Duties of beneficence are not well understood. Peter Singer has argued that the scope of beneficence should not be restricted to those who are, in some sense, near us. According to Singer, refusing to contribute to humanitarian relief efforts is *just as wrong as* refusing to rescue a child drowning before you. Most people do not seem convinced by Singer's arguments, yet no one in my view has offered a plausible justification for restricting the scope of beneficence that doesn't produce counterintuitive results elsewhere.

I offer a defense of this restricted scope by introducing the notion of *unique dependence*, a notion that is both intuitively attractive and theoretically grounded. Roughly, someone is uniquely dependent on you when you are the only individual in a position to provide assistance. Why unique dependence deserves the importance ascribed to it is further explained by the following considerations.

First, when someone is uniquely dependent on you, you are in control of her situation. But this control cannot be exchanged for a reduction of suffering elsewhere without treating the value of this person as interchangeable with the value of others. But treating a person's value in such a way does not reflect our beliefs about the dignity and autonomy of persons.

Second, the relation of unique dependence entails that one's beneficiary is *determinate*. The reason this is important becomes clear once we accept the plausible (Scanlonian) idea that moral justification depends on what determinate *individuals* could or could not reasonably accept, not on the aggregative wishes of the whole. Thus when you refuse to make the sorts of contributions to humanitarian relief efforts that Singer and others demand, there is *no* determinate individual who can claim that had you made the effort to save her life, she would almost assuredly be alive.

Addressing distant suffering is, and ought to be, the work of collective bodies, particularly those bodies that set policy. *Your* individual obligations are to those bodies. When you refuse to contribute, you are responsible for not making an effort to reduce suffering, but you are not responsible for any determinate deaths.

REASONS FOR RESCUE: AN ESSAY ON BENEFICENCE

by

Scott M. James

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2005

Advisory Committee:

Professor Christopher W. Morris, Chair
Professor Peter Carruthers
Professor Mark Graber
Professor Patricia S. Greenspan
Associate Professor Samuel Kerstein

© Copyright by

Scott M. James

2005

For my parents
Mac and Pam James

ACKNOWLEDGEMENTS

I am indebted to many people for their help. Over the years, Patricia Greenspan has been the source of much philosophical inspiration, and her constant demand for clarity and care (I hope) is reflected in the work that follows. I am especially grateful to Chris Morris and Peter Carruthers, both of whom provided detailed comments on early drafts of the thesis. Sam Kerstein, though only recently added to the committee, has been an invaluable guide in reading Kant; I profited immensely from his recent seminar on Kantian ethics. Although I can expect Sam to disagree with some of the major conclusions of the thesis, his criticisms have always been constructive and encouraging. I also want to thank Mark Graber for agreeing to serve (on rather short notice) as the Dean's Representative.

I have also received helpful comments from anonymous referees at The Canadian Journal of Philosophy and Social Theory and Practice, where parts of Chapters 3 and 6 are under review.

I benefited from a number of discussions with other philosophers and economists. Among the former, I would like to thank Mark Schroeder, David Lefkowitz, Brian Baltzly, Erich Deise, R. Jay Wallace, as well as audiences at New York University, Columbia University, and Syracuse University. Several development economists were nice enough to share their time (and patience) with me: Jonathon Tin, Jessica Priselac, and Julia Stiles.

As always, these people deserve credit for the things I got right; I am to blame for what I got wrong.

TABLE OF CONTENTS

<i>CHAPTER 1</i>	Introduction.....	1
<i>CHAPTER 2</i>	Rescue and Wrongness: A First Look at Revisionism.....	9
<i>CHAPTER 3</i>	Easy Rescue and Distant Suffering: A Closer Look at Revisionism.....	43
<i>CHAPTER 4</i>	Limiting Obligations.....	79
<i>CHAPTER 5</i>	Suffering and Practical Reason.....	113
<i>CHAPTER 6</i>	Unique Dependency.....	144
<i>CHAPTER 7</i>	Humanitarianism.....	175
<i>APPENDIX</i>	A Universe Apart.....	207
<i>BIBLIOGRAPHY</i>	216

CHAPTER ONE

Introduction

THIS IS AN ESSAY ABOUT SAVING LIVES—about those lives we are morally required to save and those we are not morally required to save. More specifically, it is an attempt to resolve a puzzle in one corner of normative moral theory. That puzzle has assumed a number of different shapes in the literature, but at its core are two deeply held intuitions. The first intuition is that you ought to prevent others from suffering needlessly when the costs are minimal. Suppose you find yourself in a situation in which someone, through no fault of his own, stands to lose something significant like his life. The intuition is that you would have a compelling moral reason to take the minimal steps necessary to prevent that loss so long as there is no comparable risk or cost to yourself. The second intuition is that morality does not require that you make a career out of saving *any* and *all* lives you can. You are not obligated, for example, to donate most of your income and time to eradicating extreme poverty abroad, despite the fact that such poverty is responsible for millions of deaths annually.

The puzzle, however, is that the first intuition implies (if not entails) the denial of the second—or so it appears. For the compelling moral reason to take the minimal steps necessary to prevent someone from losing something significant does not appear to be constrained by such things as geography, relation, race, etc. It looks to be a perfectly general imperative to prevent suffering. What is more, for any sacrifice you are required to make to prevent the suffering of one individual abroad (forgoing a nice dinner, new clothes, a vacation, photography lessons, etc.), that

sacrifice will pale in comparison to the sacrifice that individual will be required to make in order for you to enjoy that luxury. Thus there does not look to be room to endorse the first intuition without also rejecting the second.

Considerations such as these convince some that the way out of the puzzle is to dismiss the second intuition. Rather than indicating something of real moral significance, the intuition that we are not responsible for deaths abroad gives a false impression of our genuine moral responsibilities. The second intuition is distortion. This is the path Peter Singer and Peter Unger take. They have been leading a campaign to radically alter our conceptions of moral decency. In their view, it is nothing less than morally scandalous that well-off individuals like you and I do not do everything we can to reduce easily preventable suffering and death.¹ Contrary to prevailing opinion, making significant sacrifices for famine- and disease-ridden individuals abroad is not “above and beyond” the call of duty, say Singer and Unger; that *just is* our duty. Call this position *moral revisionism*.

Another way out of the puzzle is to reject the first intuition. One could retreat to the (extreme) moral libertarian view, according to which the only moral obligations agents are under are “negative” obligations, that is, obligations against acting in certain ways. On this view, you cannot wrong someone by failing to do something for her; the only way you can wrong her is by doing something to her. Quite obviously, then, there is no obstacle to endorsing the second intuition if one denies

¹ What this means in practical terms is something of an open question. While Singer and Unger urge us to contribute to emergency relief agencies like Oxfam and UNICEF, nothing in their argument requires that our efforts take this shape. If it turns out that more lives can be saved by contributing to Amnesty International, say, or corporations that invest in equity capital in poor nations, this will simply show that Singer’s and Unger’s economic, not moral, assumptions were mistaken.

that there are any positive moral obligations at all. Call this position *moral libertarianism*.²

A third way out of the puzzle is to simply deny that the puzzle is real. Despite appearances, there is no deep tension between the intuitions. This is the approach I defend in the present work. On my view, revisionism and moral libertarianism go astray in dismissing these intuitions—not because we ought to view these intuitions as sacrosanct, but because these intuitions are in fact manifestations of deeper conceptual judgments about the structure of morality. To be clear, I am not advocating the preservation of intuitions for the sake of preserving intuitions, as if pre-theoretic judgments have any better claim to the truth than theoretically-informed judgments. Rather, it is my contention that dismissing such intuitions is appropriate only if they can be shown to be the product of morally irrelevant considerations. But I submit that these intuitions are in fact the product of morally relevant—indeed, morally significant—facts about such things as wrongness and beneficence.

The key to my approach lies in distinguishing the intuitions that supposedly lead to the puzzle. According to the revisionist, there is no interesting moral difference between one's relationship to an individual involved in a local emergency, say, and one's relationship to distant sufferers. Preventable death is preventable death. But this is true only if *what explains* the imperative to rescue the individual in an emergency is the same as what explains the imperative to contribute to humanitarian relief efforts. On the revisionist approach, the first intuition is grounded in the thought that one ought to respond to the individual because doing so constitutes

² I do not consider moral libertarianism a serious competitor to the views under consideration. The only philosopher I am aware of that comes close to endorsing it is Benditt (1985). Apart from discussion of it here, I will not discuss it further.

reducing suffering or decreasing the net badness in the world or some such thing. If we accept this explanation, then it comes as little surprise that we are committed to an imperative to respond to distant suffering since doing so amounts to fulfilling one's commitment to reducing suffering, etc. But, I argue, the imperative to respond to the individual in an emergency is not, in the first instance, grounded on a reason to reduce suffering. The revisionist mischaracterizes the case. What makes your refusal to rescue the individual morally wrong is not, as the revisionist maintains, how your refusal impacts or affects some independent sum of goodness or badness. What makes your refusal wrong is how you treat *him*. There is something exceptional about your relationship to someone in an emergency situation that is absent in the case of distant suffering, viz., in the former case, but not in the latter, there is a determinate individual who has a claim on you, who is especially vulnerable to being wronged if *you* do not respond to his predicament.

This points to a more fundamental feature of morality which is obscured by the revisionist's explanation: to put it crudely, morality concerns our relationship to other persons, not to the state of world generally. The suggestion is paradigmatically non-consequentialist. What determines the moral status of our actions does not depend on the sorts of consequences we bring about, but on how we value our relationships to other persons.

The larger lesson to draw from all of this is that endorsing this explanation of the first intuition does not by itself license a rejection of the second intuition. The reason? Our relationship to distant sufferers is mediated in a way that renders unintelligible the idea that there is a determinate individual who has a claim on me *as*

stringent as the claim on me that individual in an emergency has. To be sure, the distant sufferer has *some* claim on me, but it is a different sort of claim: He can claim that my refusal to contribute lessens his chances of surviving by some tiny fraction. The individual in an emergency, however, can claim that my refusal to assist him practically guarantees his continued suffering. Insofar as morality concerns our relationship to other persons, and not to the overall state of the world, this is a difference that makes a difference.

This approach to the puzzle departs from previous approaches in several ways. Many philosophers have taken issue with the *demands* rejecting the second intuition would impose on us. There is no question that revisionism would place such demands on us, but, as I argue in **Chapter 2**, it is quite difficult to know what to make of these objections, partly because it is difficult to determine what counts as a demand in the first place. I believe that there are interesting (and novel) ways of addressing the puzzle that do not turn in any critical way on the issue of demands.

In **Chapter 2**, I argue that the wrongness of refusing to rescue individuals involved in local emergencies is not exhausted by the fact that that refusal yielded a worse outcome than other available responses. Utilitarian interpretations fail to account for several important features of emergency cases. Most notably, they fail to explain why we tend to think that someone uniquely dependent on you is especially vulnerable to being wronged by you. I present a hypothetical case in which you are forced to choose between saving a determinate individual drowning before you and sending lifesaving resources abroad. Our reaction to the case strongly supports the

idea—as yet, unexamined—that some individuals have a greater claim on our assistance than others.

This becomes all the more important when we turn, in **Chapter 3**, to the case of distant suffering and the question of whether or not distant suffering ought to be modeled on our paradigm cases. I argue that it should not. One of the distinguishing features of local emergency cases is the fact that there is a determinate individual at risk. I argue that this fact has not received the kind of theoretical attention it deserves, and I present a series of cases isolating this feature for examination. The lesson, that I go on to elaborate on in **Chapters 6 and 7**, is that only individuals—that is, determinate individuals—can make moral claims on us. This provides a sharper picture of what various individuals in need can reasonably claim from us.

In **Chapter 4**, I deviate slightly from the program by considering ways in which our obligations to those in need might reasonably be limited. Standard objections to revisionism maintain that revisionism cannot account for reasonable limits on what we can be required to do. The literature contains two general approaches to the formulation of limits: the personal point of view approach and the impersonal point of view approach. I argue that neither approach is conclusive, though I believe that Murphy's (2000) emphasis on compliance and the way in which noncompliance places unfair burdens on optimal compliers can and ought to be exploited. I do so in **Chapter 7**.

In **Chapter 5**, I return to the question of our reasons to respond to suffering, arguing that there is no standing reason to reduce suffering. We need not accept the claim that there is a standing reason to promote the good, nor must we accept the

claim that the badness of suffering is itself a reason to reduce (or prevent) it. This leads me to question one of the central assumptions underlying our discussion, viz., that suffering is unconditionally bad, that its badness is not conditioned on anyone's having a reason for reducing it. In its place, I argue that suffering is bad, if it is, in virtue of the badness *to* sufferers. In short, our reason to respond to the serious suffering of others has as its object sufferers, not suffering, for it is persons and their reasons that matter and not their experiences per se.

Chapters 6 and 7 represent the heart of my view, for it is here that I lay out the moral grounds for rescue: unique dependency. This is the idea that you are the only person in a position to provide the necessary assistance to an individual in an emergency situation. As I argue, you have a greater responsibility for those who are uniquely dependent on you than those who are not so dependent on you. Your moral relationship to someone uniquely dependent on you is characterized by the expectation that you will not see the value of assisting him as interchangeable with the value of assisting others. There is a second-personal aspect to unique dependency that sets it apart from humanitarianism in general. And this is captured by the thought that moral agents display *inter alia* a general and abiding reluctance even to consider exchanging the life of someone uniquely dependent on them for other lives elsewhere because this is part of the normative standard for being a moral agent. I draw on two sources as support: Kant's notion of dignity and Joseph Raz's notion of incommensurability, both of which underline the relational aspect of morality. What is more, contemporary contractualist frameworks, such as Scanlon's (1998), offer a

plausible *realization* of the idea that morality is essentially a matter of how we relate to others.

In **Chapter 7**, I argue that unique dependency can be extended quite naturally to cover cases in which there are multiple donors. Suppose five of you happen upon a drowning child. While the drowning child is not uniquely dependent on *you*, he is uniquely dependent on the group of you, since the group's failure to mount a response practically guarantees the child's imminent death. Second, this account can also be extended in such a way as to illuminate the puzzle of humanitarian relief: it can be argued that while individuals living in extreme poverty are not uniquely dependent on you or me, they are uniquely dependent on groups of which you and I are members, for example, a wealthy nation like the U.S.

The **Appendix** offers an empirical aside on the face of extreme poverty in our world today. It is meant to emphasize the sense that we as group can be thought of as the single entity in a position to rescue millions of others, for that is very nearly the case.

CHAPTER TWO

Rescue and Wrongness: A First Look at Revisionism

REVISIONISM IS A CLAIM about the moral requirements to prevent suffering. According to the view, we have a conclusive reason to prevent as much innocent suffering as we possibly can without regard to where that suffering occurs or how it was caused. Refusing to do so constitutes a serious moral failing. Revisionism thus embodies a basic challenge to commonsense morality: it challenges the idea that while we are not as charitable as we could be, most of us can nevertheless claim to be morally decent people. If revisionism is correct, this idea must be rejected. For it is an immediate consequence of this view that we are required to make large sacrifices in order to bring ourselves in line with moral decency, something almost none of us does. In this and the next chapter, I will argue that this view is mistaken. In Chapter 7, I will propose an alternative account of our moral requirements to distant sufferers.

The most popular objection to revisionism has centered on the significant demands it places on complying agents. It is argued that because no normative moral theory ought to be accepted that requires complying agents to make substantial and sustained sacrifices, we ought not accept revisionism. While I do not doubt that there is an intuitive sense in which revisionism requires complying agents to make unreasonable sacrifices, I am not convinced that this fact (if it is a fact) can be deployed as part of a non-question-begging objection to revisionism. It seems to me that revisionism is flawed on a more basic level, viz., at the level of specifying the content of our reason to render aid to others in need.

The main line of argument for revisionism, and the one I'll be considering in this essay, goes back to Singer (1972): it involves drawing parallels to paradigm cases of easy rescue—or what I'll sometimes refer to as local emergencies. The argument, in its simplest form, is that our reason to prevent distant suffering derives from the same source as our reason to prevent local emergencies. The content of this source varies slightly across the different authors we'll be confronting, but the general idea is that our reason to prevent suffering derives from a more general reason to prevent something bad—its corollary being a reason to promote the good.

After offering preliminary reasons for doubting the coherency of the over-demandingness charge, I set out in Section II a general argument schema as a way of capturing the revisionist argument. In Section III, I consider several broadly consequentialist interpretations of our standing reason to render aid. I argue that the wrongness of refusing to rescue individuals involved in local emergencies is not exhausted by the fact that that refusal yielded a worse outcome than other available responses. The consequentialist interpretations fail to account for several important features of those cases. This becomes all the more important when we turn, in Section IV, to the case of distant suffering and the question of whether or not distant suffering ought to be modeled on our paradigm cases. I argue that it should not.

Throughout our discussion, however, the reader should bear in mind that rejecting revisionism does not warrant the claim that we have no reason to render aid to distant sufferers. I will argue in Chapter 7 that we have such a reason. What I deny is that this reason flows from the same source as our obligations to local individuals.

I. Doubts about “Over-Demandingness”

Much of the discussion of revisionism has focused on the question of demands. Arguably, the concern has been around since Sidgwick, though it is was perhaps Bernard Williams (1973) who first articulated the objection that has come to be known as “over-demandingness.” According to Williams, because utilitarianism places unreasonable demands on complying agents, such that they are unable to establish and pursue meaningful personal projects, utilitarianism ought to be rejected as a normative guide to moral action.³ But Williams’s argument would seem to apply to all versions of revisionism—be they consequentialist, deontological, perfectionist, or whatever. This should be encouraging to critics of revisionism, for if a Williams-style objection could be shown to stand against all versions of revisionism, it would seem that we could set aside revisionism in one fell swoop. Unfortunately, the over-demandingness claim is quite difficult to pin down.⁴

First, among the things an assessment of demands will require is a baseline. For without establishing some point from which losses can be measured, it will remain unclear when in fact someone incurs losses through her actions (or inactions).⁵ But how to establish an acceptable baseline is far from clear. We could opt for a baseline of the factual status quo, according to which “a moral theory makes demands on an agent to the extent of any losses he does or would suffer *just because* of his

³ See Chapter 4 for an extended discussion of Williams on the personal point of view.

⁴ The following points have their origin in Murphy (2000), Chapter 1-2.

⁵ Another problem also lurks in the wings: Is it even possible to make interpersonal value comparisons? That is, are the costs to me of doing Φ commensurable with the costs to you of doing Φ ? See, e.g., Otteson, (2000).

compliance with the moral theory.”⁶ So if Jean and Kelly each had to give up a thousand dollars in order to discharge their moral obligations, and if doing so would leave them at roughly the same level of well-being, then, according to the baseline of the factual status quo, the costs to Jean and Kelly are the roughly same.

But if Jean had to return the money to its rightful owner—because Jean stole it—while Kelly had to send the money to humanitarian relief efforts, the costs do not appear the same at all. The reason is plain: Jean is not *entitled* to the money in the first place. Returning the money, while perhaps regrettable from Jean’s perspective, should not count as a cost to Jean, at least in the same sense that Kelly’s giving his money to humanitarian relief is a cost to Kelly. But now we’ve exchanged a factual or descriptive account of the baseline with a normative account, according to which a moral theory places demands on complying agents only to the extent that they are required to give up something they are otherwise entitled to. But now we have need for a conception of possessory entitlement, and this lands us in the problem of having to frame a non-question-begging account of possessory entitlement.

According to the utilitarian, we are entitled only to those things we are not required to bestow on others. But of course the utilitarian argues that we are required to bestow all we have onto others, short of endangering ourselves; hence, if we determine the baseline according to utilitarian standards of entitlements, then the problem of over-demandingness dissolves altogether. As such, it is no objection at all to say that utilitarianism—or, for that matter, any version of revisionism—is over-demanding. For the sacrifices you make for the sake of others are not over-

⁶ Murphy (2000), p. 35.

demanding since you were never entitled to those goods in the first place.⁷ Something has gone wrong.

But the problem goes deeper still. Recall that the charge of over-demandingness was meant to function as a theoretical constraint: any moral theory that made excessive demands on complying agents was, for that reason, objectionable. But if the characterization of our baseline must be made from within a given moral theory—that is, according to the standards provided by that theory—then we cannot pretend that the charge of over-demandingness carries force against other theories, for those theories will not share the same normative characterization of the baseline.

We might also raise doubts about the cogency of the claim that a moral theory ought to be rejected if it makes excessive demands on complying agents. If we set aside the problems we just considered and suppose that a theory indeed requires significant sacrifices of complying agents, should we count this as a sufficient reason to reject that theory? Perhaps not. For we must not forget that demands are highly relative to the circumstances within which one finds oneself. For example, most of us in the affluent West would hardly say that refraining from stealing a hundred dollars (under everyday circumstances) places a high demand on us. But of course in many places of the world, such restraint would be very costly to the agent, for example, when that money could easily be stolen from corrupt officials, say. Or imagine a situation in which, as a government official, you are required to torture a political prisoner in order to save not only your own life but the lives of your family members.

⁷ For an account of possessory entitlements at the other extreme, see Nozick (1974).

Most of the commonsense prohibitions we abide by in the West (prohibitions against lying, stealing, murdering, etc.) rarely strike us as costly. But this of course is a consequence of the fact that typically we do not stand to gain substantially by such acts. What is more, stiff legal sanctions further reduce the advantage of violating these commonsense prohibitions. Consider, however, those places in the world in which individuals do stand to gain substantially from such acts. In places where the rule of law is broken or simply nonexistent, complying with moral prohibitions not to steal, lie, or worse can be as demanding as any of the active (or positive) duties we in the West judge to be demanding. Do we want to say, therefore, that we should reject theories that prohibit such things as murder and stealing? If it is true that a moral theory ought to be rejected if it makes excessive demands on complying agents, then many commonsense prohibitions face rejection as well, whatever their moral foundations.

To address this concern, the defender of the over-demandingness charge, has to say that there is something particular about the demands of beneficence or positive duties that make them objectionable. For example, it might be asserted that it is only when a theory makes *active* (as opposed to merely passive) demands on complying agents that it is objectionable. Murphy, for example, thinks that if the over-demandingness charge is to have any hope of sticking, it must address the distinctive demands of positive duties. This is a step in the right direction, though I believe that Murphy's account must ultimately be rejected (see Chapter 4). For time being, I believe we should set aside the question of over-demandingness, at least until we have a better sense of what sorts of demands beneficence is supposed to entail.

II. Reasons to Render Aid: A Revisionist Schema

As a way of framing our discussion, I want to introduce a general argument schema. Quite deliberately, I have left out much of the important content. The point is to establish, on behalf of the revisionist, a common framework within which most revisionists operate.⁸ I'll first offer a rough sketch of the schema so the reader gets a sense of its various steps and its general shape, before going on to unpack each of the steps. Finally, I will raise doubts not only about several of the steps, but, more importantly, about the wisdom of the schema itself.

The schema is a piece of moral reasoning; the conclusion is that I have a conclusive reason to render aid in these circumstances. As a piece of moral reasoning, it depends on answers to questions at two levels. On the theoretical level, it depends on the answer to the question, what is the content of our standing reason to render aid in cases of easy rescue, assuming there is such a reason?⁹ On the practical level, it depends on the answer to the question, does this occasion count as an instance of easy rescue? Together, these answers provide content to our schema which in turn should determine, for each set of circumstances C in which someone S find herself, whether or not S has a conclusive reason to render aid in C . Here then is an outline of the revisionist schema.

The first step is to identify a set of circumstances, usually hypothetical, in which we agree that anyone in circumstances of that sort ought to render aid. It is

⁸ Among the revisionists, I count Peter Singer (1972, 1993), Peter Unger (1995); Shelley Kagan (1989); and Garrett Cullity (1994). Some are more explicit than others about how our practices ought to be revised.

⁹ Chapter 5 takes up the question of a standing reason for beneficence.

argued next that what best explains this agreement is that we recognize, first, that these cases are all instances of a type of situation T and, second, that we have a reason to render aid in situations of type T. So if I recognize that situation t is a situation of type T, then I have a reason to render aid in t . But we also see that this reason can sometimes be outweighed by more compelling considerations. I might recognize that I have a reason to render aid in t but that I have other, more stringent, reasons permitting me to withhold aid in t . This may include, for example, having a second-order reason not to take certain first-order reasons as conclusive. So it is allowed that if I do not have a more stringent reason not to render aid in t (weightier than the reason to render aid in t , that is), then my reason to render aid in t becomes a conclusive reason to render aid. This much should be common ground between the revisionist and non-revisionist. It is the next step that divides the two.

Recall that the revisionist wants to argue that I have a conclusive reason to render aid in an extensive range of cases, cases in which I do not commonly think of myself as having a conclusive reason to render aid. The case of distant suffering is a case in point. To show this, however, the revisionist needs to show that in these more controversial cases I do not have a more stringent reason not to render aid. But the revisionist needs to be explicit about the content of our reason to render aid. As I suggested above, the standard revisionist line maintains that this standing reason derives from a more general reason to promote the good. If our standing reason indeed flows from a reason to promote the good, then unless I can show that not rendering aid in these circumstances will promote more good than rendering aid, then

my reason to render aid becomes conclusive. Let's look more carefully at each step of the schema.

The first move in the schema is abductive. The aim is to derive a reason to render aid in a set of circumstances from the data of our moral responses. The revisionist reasons that since you judge that it would be seriously wrong to do nothing to prevent someone from suffering a significant loss in circumstances *c*, in circumstances *d*, and in circumstances *e*, and since in *c*, *d*, and *e*, what matters (morally) is that the potential aid-donor could have easily prevented the loss by sacrificing only a minor good to himself, it is supposed to follow (abductively) that our judgments regarding *c*, *d*, and *e* flow from a more general judgment according to which it is always wrong to do nothing to prevent someone from suffering a significant loss under any circumstances in which preventing such a loss can be had by making a minor sacrifice to oneself.¹⁰ To account for the pattern of case-specific judgments we tend to make, a covering principle (or what I call a "standing reason") is posited that best explains this pattern. This is part of what Kagan (1989) calls the "drive to explanatory coherence," and here the pressure to posit general principles of beneficence becomes strong. Two things, then, call out for explanation: first, the data from which the revisionist hypothesizes and, second, the general features that constitute our circumstance type.

First, what sorts of cases does the revisionist draw on? The case of the child drowning in a shallow pond a few feet away has, by now, become the trusted "go-to"

¹⁰ Cf. Singer's (1972: 231) loose formulation: "if it is in my power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it."

case for revisionists.¹¹ Confronted with someone *S* who refused to rescue the child and thus allowed the child to drown, most of us judge that apart from being a cruel and heartless individual, *S* *wronged* that child. (Some are tempted to say that in these circumstances the child had a right to be saved by *S*, given how easy it would have been for *S* to do so, but I want to leave the issue of rights to one side.) Unger introduces the case of the severely injured bicyclist needing a ride to the hospital in your vintage sedan.¹² In this case, too, refusing to give the bicyclist a ride to the hospital on the grounds that he will ruin your car's upholstery is considered by most to be morally repugnant. Unger contends that you wrong the bicyclist by refusing to transport him to the hospital. And when he loses his leg, most people (supposedly) believe he has a legitimate grievance against you. Likewise with respect to Hershenov's lifeboat case: When you refuse a man-overboard a space in your lifeboat because it will mean forever parting with a large chest filled with clothes, and he subsequently drowns, you not only exhibit a shameful lack of benevolence, but you wrong him.¹³

Suppose that we do find the failure to act in these cases morally wrong (and not merely cruel, insensitive and the like).¹⁴ What explains this? Why is it that most of us are prepared to censure someone morally for failing to act in these

¹¹ The example originated with Singer (1972) but turns up in nearly every related discussion, the present work not excluded.

¹² Unger (1995), p. 24 ff.

¹³ Hershenov, (XXXX). I have modified Hershenov's original example slightly. Particular cases need not be fanciful: take the now-famous case of Kitty Genovese, who was brutally raped and murdered outside her apartment building while more than two dozen residents of the apartment, despite hearing her cries for help, stood by and did nothing. The outcry to bring legal action against them suggests that they wronged her. See Feinberg (1984).

¹⁴ To establish the existence of moral requirements, we must do more than point out the character flaws that presumably attend the failure to render aid in these circumstances. For there is no inconsistency, at least at this point, in claiming that Smith is cruel and heartless for ignoring the drowning child and yet insisting that the child has no right to Smith's assistance, that Smith has no duty to aid the child.

circumstances? There are two ways of reading this question, only one of which will concern us. The first reading, which I'll leave to one side, asks the psychological question: which facts, as a matter of psychology, were responsible for eliciting the response they did? The second, and pertinent, reading asks the normative question: which facts justify the response most of us are disposed to make? Why, that is, would it be wrong (assuming it is) for someone to refuse to render aid in those circumstances? The commonsense response seems to be something to the effect that, 'You could have prevented serious harm to someone at a minimal cost to yourself and you did nothing.'

Before we can spell out the details of this response, we should get clear on what features all of these cases are supposed to have in common. We need to articulate, that is, a situation type. For unless we can group these cases under a general situation type, it will remain unclear whether or not a given situation counts as an instance of that type and, hence, whether or not we have a reason to render aid in that instance.

Consider, then, the following provisional list as constituting the relevant type:

- a) an innocent person stands to suffer a significant loss, like a limb or his life;
- b) some agent *S* is in a position to prevent such a loss, either through her actions or her resources;
- c) *S*'s executing those actions or contributing those resources can be expected to result in only a trivial cost to *S* when compared to the loss the victim faces;
- d) *S*'s efforts at rescue under the circumstances do not require *S* to perform actions that are prohibited by commonsense morality; and
- e) *S*'s efforts at rescue are not likely to increase the number of subsequent individuals needing rescue.¹⁵

¹⁵ On some accounts, (d) will simply rule out (e), since some will argue that there is a commonsense prohibition against acting in ways that will worsen the circumstances in the long run, even if in the short term they are improved.

The revisionist argues that this list of features constitutes a situation type in response to which most of us judge that aid ought to be rendered. I will call this type *ER* for “easy rescue.” Note, first, that this list appears to describe the cases we considered above. If this is right, then those cases are just different instances of *ER*. The drowning child, the injured bicyclist, Ms. Genovese, and so on are all in imminent danger of suffering a significant loss. And, *ex hypothesi*, at least one individual happens to be in a position to prevent such a loss by incurring a minor cost. Furthermore, nothing in the examples suggests that preventing that loss would be morally prohibited. We might conclude, then, that it is the combination of these features, together with some reason to render aid in the presence of these features that best explains why we tend to level harsh moral criticism at individuals who fail to act in relevantly similar situations.

To summarize, we began by citing a handful of cases that all seemed to elicit a similar response. We then explained this response by appeal to *ER* and a reason to render aid in *ER*. This sets up the next step in the schema, which involves introducing a novel case and arguing that this case, too, is an instance of *ER*.

Recall that part of the revisionist claim is that the moral requirements to rescue others in need are extensive; that is, they are not confined to what might be called local emergencies. Here is where the revisionist makes good on that claim. For the revisionist is going to argue *inter alia* that the case of distant suffering is an instance of *ER*. According to the revisionist, we stand in similar moral relationships to the famine- and disease-ridden individuals abroad as we do to the drowning child,

the injured bicyclist, and so on. If the revisionist is right, then it follows that you and I have a reason to render aid to those individuals (through the relevant channels, presumably), since you and I have a standing reason to render aid in any situation of the designated type.

We will return in the next chapter to the claim that distant suffering counts as an instance of easy rescue. I will suggest several reasons why it should not be. In the meantime, our task is to elaborate on how this standing reason might be analyzed. For the only way to know if the differences that separate the case of distant suffering from our original (local emergency) cases make a moral difference is if we already have a conception of morality in mind, one that offered a procedure to determine how these differences ought to be weighed. A decisive answer to the question, is distant suffering a case of *ER*, cannot proceed independently of a moral theory that assigns weights to the various reasons for and against responding to others in need.¹⁶ Indeed, much of the present work is aimed at supplying such assignments.

III. Consequentialist Reasons to Render Aid

Our concern in this section is to examine the content of this so-called standing reason to render aid.¹⁷ In particular, I want to evaluate a family of consequentialist interpretations beginning with classical utilitarianism. As we'll see, it is no coincidence that most revisionists tend to espouse a consequentialist conception of morality. Of course a complete evaluation of this reading would require an extensive

¹⁶ Unger seems to deny this, since he tackles the question without appeal to any basic normative assumptions other than what he calls "our moral common sense." See Unger (1995) Chapter 7.

¹⁷ Again, I am assuming that there is such a standing reason. In the end, I do not think there is such a reason (see Chapter 5); however, the arguments of the present section are more modest insofar as they merely challenge the content of those reasons.

evaluation of consequentialism itself, which I cannot hope to do here. For our purposes, it is enough to consider how various consequentialist stories can be deployed in understanding the wrongness of refusing to render aid in cases of easy rescue. Roughly put, consequentialism contends that the wrongness consists in failing to bring about good consequences. Classical utilitarianism, as a species of consequentialism, contends that preventing bad consequences is apiece with maximizing utility.

These explanations, I submit, are deficient on several fronts. In the context, say, of the drowning child case, the wrongness of ignoring the child is not exhausted by the fact that one's omission resulted in a worse overall state of affairs than other available responses. First, the classical utilitarian interpretation cannot account for the sense that *that* child has (or would have) a legitimate grievance against *me* for failing to aid him, a grievance that is not mediated by how things are assessed from an impartial standpoint.¹⁸ Second, I will argue that promoting the good does not adequately explain why, as the donor, I have a special obligation to *justify* my particular act of omission.

Recall that the revisionist schema begins by adducing a set of judgments concerning easy rescue, from which the revisionist draws out the standing reason that best explains these judgments. Again, this is not the task of identifying what, as a matter of fact, elicits these judgments; rather, it is the task of justifying these judgments by reference to a more general reason that one accepts. Let's begin with Kagan's (1989: 17) suggestion that "the best explanation of the various judgments

¹⁸ In Chapter 5, I challenge the more fundamental assumption that there is a fact of the matter as to how things are "from the impartial standpoint," that we can, in other words, make judgments about the "overall state of affairs."

that the moderate wants to make involves the acceptance of a standing, pro tanto reason to promote the good.”¹⁹

Within classical utilitarianism more generally one finds varying interpretations of the notion of “promoting the good.” Rather than canvassing and critiquing these various interpretations, I want to propose a neutral interpretation, one that leaves it open how value is ultimately supposed to be understood. To wit, to have a reason to promote the good is to have a reason to bring about the highest net gain in good states of affairs over bad, where “good states of affairs” is understood to mean states of affairs in which as much good as possible is realized, whatever “the good” turns out to be.²⁰ In situations in which no net gain in the good is possible, then one has a reason to secure the lowest net loss in good states of affairs as possible.

The reason for speaking of “net gains and losses” has to do with the need to account for the losses someone may accrue through her responses, under which I include both commissions and omissions. Suppose, for example, that the expected good of some action *A* is less than the cost to me of performing *A*. There would thus be a net loss in performing *A*. The world, as it were, would be a worse-off place were I to perform *A*. Utilitarianism would therefore forbid me from performing *A* under the circumstances.

¹⁹ Kagan’s aim is actually two-fold: on the one hand, he wants to show that morality does not allow agents the *option* or prerogative not to maximize the good. It is this aspect of Kagan’s view—the defense of stringent positive duties, that is—that will concern us here. But the other side of the coin, according to Kagan, is that morality does not contain *constraints* on what we might do to each other in the service of promoting the good. We will have occasion to address this side of Kagan’s view, though it will not bear directly on our discussion.

²⁰ To be sure, one thing “the good” cannot turn out to be is moral good, since morality or rightness, according to classical utilitarianism, is defined *in terms of* the good, so one needs an account of the good that is wholly independent of moral concepts to avoid circularity.

By contrast, if the expected good of some action *B* is substantial (e.g. preventing a child from drowning) in relation to the cost to me of performing *B* (e.g. a ruined pair of shoes and the physical effort of wading into the pond), and if there is no other response that can be expected to yield a lower net loss in good states of affairs, then morality requires that I perform *B*. On this view, the only justification I can have for not rescuing the child is that I can expect a lower net loss from not rescuing the child.²¹

Importantly, appeal to expected net gains and losses in overall states of affairs exhausts the content of our standing reason to render aid in a given set of circumstances, according to the classical utilitarian. While you may have other reasons for rescuing the child, the only *morally relevant* reason is the fact that rescuing the child will prevent a substantial net loss in good states of affairs; refusing to rescue the child is wrong only because, from an impartial standpoint, you allowed a sizable net loss in good states of affairs.

There are several advantages of thinking about promoting the good in terms of net gains and losses. First, as mentioned above, appeal to net gains allows us to explain why, in some circumstances, we are not required to aid others. When we subtract the expected badness of performing *A* from the expected goodness of performing *A* and discover that, in these circumstances, the net goodness is much lower than other available responses—because, for example, the harm I sustain in aiding another is much greater than the benefit I bestow on him—we can account for

²¹ As an example of such a case, imagine that I stand midway between two ponds. While our one child cries out for help in the East Pond, three children call out for help in the West Pond. I save the three, allowing one to die. From an impartial standpoint, my action resulted in a lower net loss than what we could have expected had I decided to save only the one. This fact is both a necessary and sufficient condition for justifying my action, according to the classical utilitarian.

the thought that there are limits to what we are required to do for others. After all, as Mill (1863/2001: 17) noted, “a sacrifice which does not increase, or tend to increase, the sum total of happiness, [utilitarianism] considers wasted.”

The other side of the same coin is the idea that net gains (and losses) can explain the wrongness of not rendering aid in many of our original cases. Most of us have the intuition that a large part of what makes failing to rescue the drowning child wrong is the fact that it would cost so little. The classical utilitarian captures this idea in terms of net gains and losses. When we subtract the badness to me (but not exclusively to me) of rescuing the child from the goodness to the child (but not exclusively to the child) of being rescued, we are left with a sizable net gain in good states of affairs over bad, since the overall goodness consequent to the rescue is diminished only slightly by the overall badness of performing the rescue. This can easily explain why we find “easy rescue” morally compelling.

Incidentally, it is easy to see how this interpretation of our standing reason to render aid leads to revisionism. Recall that it is a consequence of revisionism that the moral requirements to render aid to others are demanding (though, for reasons discussed above, not necessarily *over*-demanding) for individuals like you and I. To see this, consider occasions on which I am not required to rescue another in need. While it is true that I was not required to perform *A* above, this was only because the resulting overall state of affairs would have been lower than alternative responses (e.g., doing nothing). This is not the same thing as saying that I am not required to perform *A* because the resulting overall state of affairs would be especially bad for

me, that it would be particularly demanding on *me*.²² Just because a given act of rescue will result in substantial badness to the donor does not necessarily mean that the donor does not have a conclusive reason to render that aid. Chances are, she does.

It is only when the badness to the donor—offset by the goodness to the recipient—results in a net loss (or lower net gain) in good states of affairs by comparison to alternative responses that the donor is relieved of her duty to render aid. But this leaves open a considerable range of cases, cases in which rendering aid will be very costly to me. We can imagine a scenario in which some action Φ can be expected to result in substantial badness to the donor (e.g. losing a leg); however, because the net loss in good states of affairs (e.g. preventing someone from losing her life) is expected to be lower for Φ than for any alternative response, the donor is required under the circumstances to perform Φ .

But even if promoting the good is sometimes not demanding, it is almost always extensive. For with only a little imagination, we can conceive of any number of activities that we could be doing at this very minute, whose outcomes would be of higher net gains in good states of affairs than our present activity. Whether it's earning a second income to send abroad or working at a local shelter or lobbying for human rights or open markets, there is no perceivable limit to the extent of beneficent activities one could be engaged in at any given minute. When “the requirement to pursue the good pervades an agent's life—all its aspects, every moment,” it is small wonder that our duties to render aid are extensive.²³

²² This would be an agent-relative reason why I am not required to rescue the child, something the classical utilitarian would reject.

²³ Kagan (1989), p. 19.

IV. Rescue and Wrongness

In critiquing this interpretation of our standing reason to render aid, I want to restrict my focus to how the classical utilitarian handles our original cases, viz., the drowning child, the injured bicyclist and so on. I want to frame my critique around the following question: does the classical utilitarian provide an adequate account of the *wrongness* of refusing to rescue the imperiled individual? The answer, I argue, is no.²⁴ In rough form the answer goes like this. On the classical utilitarian view, refusing to rescue the imperiled individual is wrong, if it is, only because of the way in which such a refusal impacts the overall state of affairs. But this fails to make sense of the observation that you have directly wronged that individual. The circumstances of the situation are such that you stand in a special relationship to him, one that makes him especially vulnerable to your inaction, but this is something the utilitarian explanation fails to capture.

According to the classical utilitarian, what makes refusing to aid the imperiled individual (morally) wrong is that your omission resulted in a substantial net loss in overall good states of affairs, which is not offset by the loss you would have had to sustain had you attempted a rescue. What legitimizes our moral condemnation of you is that you refused to prevent an increase in bad states of affairs over good. I do not wish to quarrel with the claim that refusing to rescue the imperiled individual results in a net loss in good states of affairs: the world would seem to be a worse place for

²⁴ In the next chapter, I expand in this reason, linking to other reasons.

your omission.²⁵ What I deny is that this adequately explains the moral wrongness of refusing to render aid. Among the things missing from this account is an explanation of why your omission should be particularly egregious from the point of view of the imperiled individual. According to the classical utilitarian, the imperiled individual has a complaint against you *only derivatively*, that is, only by reference to how things look from an impartial standpoint. So when the injured bicyclist expresses moral outrage at you (for his having had to lose his leg because you refused him a ride to the hospital), the only morally relevant complaint he can have against you, according to the classical utilitarian, is that from the point view of the universe you failed to prevent a substantial loss in good states of affairs by refusing him a ride. The bicyclist is restricted to the following complaint: My loss represents a loss to the whole, and it is because the value of the whole is lessened that your action was wrong.

But the content of the imperiled individual's complaint goes beyond this. For surely the injured bicyclist's complaint involves a claim about how you treated *him*—quite apart from the cumulative effect of his subsequent loss. The wrong you commit in refusing to assist the bicyclist has everything to do with the very special relationship that you bear to that individual, a relationship that is not mediated by how your response will contribute to the subsequent distribution of good and bad states of affairs. To see this, consider the following hypothetical scenario.²⁶

²⁵ This is not to say that, in the end, I find talk of “overall good or bad states of affairs” meaningful. In Chapter 5 I offer an alternative account of value which has the effect of denying that such notions have meaning.

²⁶ This example is adapted from Gomberg (2002).

As the owner of a very expensive pair of boots, you decide one day to sell the boots for a handsome price and donate the money to relief efforts abroad. You don the boots and strike out in the direction of your local pawnshop. On the way there, however, you come across a child drowning in a shallow pool. But before rushing to the child's aid, you confront two facts: first, you need the assistance of another person in order to remove the boots(you've nearly outgrown them), but no such assistance is available; second, if you enter the pool in the boots, they will be ruined and, once ruined, worthless. You reason that the money you planned to send abroad could save a few more lives than the single life you could save here and now. On those grounds, you allow the child to drown and continue on to the pawnshop where you sell the boots for humanitarian ends. According to Singer, you did the right thing. It seems clear, too, that Unger and Kagan would have to accept Singer's judgment: selling the boots would yield a smaller net loss in good states of affairs than rescuing the drowning child.²⁷

But is this really what morality requires? Are you really obligated to sacrifice that child's life for the sake of several lives abroad? To be convinced of such a thing is to be in the grips of the idea that net losses and gains exhaust the content of our moral judgments. But it seems clear that there are considerations other than overall utility rankings that count in favor of rescuing the child.

It seems clear that you wrong the child by not rescuing him. But what makes your refusal to rescue him wrong? In these circumstances, you wrong the child in part because, to put it crudely, you put the interests of the whole ahead of the interests

²⁷ A notable utilitarian exception might be R.M. Hare (1981), whose two-level approach to moral thinking might be able to account for the intuition that you ought to rescue the drowning child. I discuss and ultimately reject Hare's approach in Chapter 4.

of that child. To put the interests of the whole ahead of the interests of this child, you have to disregard certain features of your relationship to that child—first and foremost, that he and his relations rely on *you* to avert his tragic death. For, in this situation, it is you and only you who can prevent his death. When you disregard his pleas, you are disregarding a burden that falls—if it falls anywhere—on you. It is you whom the circumstances have picked out. It would thus be morally irresponsible of you to disregard the child since you bear a special relationship to him, *even though* your disregard (and subsequent humanitarian efforts) is likely to yield a lower net loss in good states of affairs.²⁸ Again, the classical utilitarian will argue that such disregard is morally required of us. But this is to assume that such disregard does not count as a serious moral failure, which it clearly seems to be.

In order to justify your omission, you must see the needy individuals abroad as i) having an equal claim to your assistance as that child and ii) constituting a greater potential loss in good states of affairs, viewed impartially, than the loss of this one child. As I've already mentioned, I do not wish to take issue with (ii)—at least not here.²⁹ The question is whether other considerations, considerations that do not make reference to overall states of affairs, ought to bear on your moral deliberations. I think the answer is yes. Part of my reason involves the denial of (i).

In order to assume that the individuals abroad have an equal claim to your assistance, it must be the case that they are the direct potential beneficiaries of your assistance. But if it is true, as the revisionist thinks, that for any distant person *X* under life-threatening conditions, *X* has an equal claim on your assistance as the

²⁸ I have attempted to work out the details behind this claim in my “Good Samaritans, Good Humanitarians: The Case for Unique Dependence,” under review.

²⁹ See Chapter 5.

drowning child, then it must also be true that *X* can reasonably expect to be assisted by you. But it is false that *X* can claim such a thing, for the intermediation of a relief agency upsets any direct assistance.

From the point of view of *X*, it is false to say, as *X* might, ‘I have a unique right to your assistance.’ For unless *X* has a reason to think that you are in a unique position to avert *X*’s impending death, the most *X* can claim is that you have an obligation to give to relief agencies, whose actions might avert *X*’s impending death. (But, if I’m right, then your obligation, if any, to give to relief agencies cannot rest on the fact that any one individual’s life will be saved by my contribution.³⁰) As we’ll see in the next section, your refusal to contribute does not guarantee *X*’s imminent death, nor does your contribution guarantee that *X* will live. While it is true that *X*’s life does, in the present circumstances, depend on outside assistance, it does not follow that *X* can legitimately rely on any one individual’s beneficence so long as beneficence consists, in this instance, of contributing to humanitarian relief efforts.³¹

But the drowning child is another story altogether. For unlike the potential beneficiaries abroad, your failure to assist the drowning child guarantees his death. His fate is in your hands. To disregard this fact is to disregard a unique feature of your relationship to that child. Moreover, such a fact places a special burden on you, a burden unlike the one distant suffering (putatively) places on you. For there, no one individual has a special claim to your assistance. But not only do you stand in a special relationship to that child, but you stand in a special relationship to the

³⁰ Some of the implications of this view are explored in my “Benefits without Beneficiaries: Notes on Moral Indeterminacy,” under review.

³¹ In the language of rights, it is commonly assumed that rights are held against certain persons. So the revisionist shoulders the added burden of explaining how, for some individual *A*, *A* can have a right to be aided by no one in particular. See Raz (1986), p. 180 ff.

community within which your actions (or inactions) take place. After all, one feature of morality I'll explore later is that we come to have expectations about the sorts of things other people ought to see as reasons. We rely on others to share our commitment to certain reasons for actions, and such reliance forms a separate category of moral reasons. Let me try to spell this out.

A large part of what makes your omission in the drowning child case wrong is the fact that you have no legitimate *reason* for your omission; as such, you cannot justify your omission to those that are subsequently harmed. After all, the parents of the drowned child will want to know, first and foremost, what *reason* you had for not stooping down and scooping their child from the pool. The fact that you had no reason is a distinct—and, I want to argue, morally significant—fact that goes beyond the fact that your omission yielded a net loss in good states of affairs.³² It reflects a flagrant disregard for the reasonable expectations of others, among them, the child's parents. The classical utilitarian will argue that, of course you had no reason, since your omission could be expected to result in a net loss. But this misses the point.

For the parents of the child are asking you to account for your omission, to bring it into line with reasons they could understand, if not accept. Why? Because they could have reasonably expected that someone in your position would have had a compelling reason to rescue their child.³³ Indeed, we operate under such expectations all the time. We assume others can see a reason not to lie to us or steal from us or assault us; correlatively, we assume others can see a reason to take easy steps to

³² The details of this conception of wrongness and its implications for beneficence are spelled out in Chapter 6.

³³ This does not mean the parents must be in possession of such expectations. It is enough that it would be reasonable for anyone to possess such expectations, thus giving anyone a reason to object to your omission.

prevent a serious loss, and these are not grounded solely, if at all, in reasons to promote the good. These interpersonal expectations thus form a separate category of reasons—separate, that is, from reasons involving only reference to expected outcomes. In short, when we act contrary to these expectations, we are called upon to provide some justification for our (in)action. According to the classical utilitarian, there is no separate realm of justification, at least one that does not in the end contribute to the overall distribution of good and bad states of affairs. But this then leaves unaccounted for the sense that failing to assist the drowning child or the injured bicyclist calls for a specific sort of justification, one that is owed directly to those individuals (or their benefactors). It is specific because it picks out specific individuals. According to the classical utilitarian, there is no independent need to justify one's omission unless doing so may contribute to a net increase in good states of affairs over bad. And even then, given the strictly impartial evaluation of your omission, the justification need not be directed at anyone in particular.

V. Source, Focus, and Indirection

The problem I have identified might be recast in terms of *source* and *focus*. According to classical utilitarianism, the source of the rightness or wrongness of your action lies in the net gains and losses in good states of affairs. This would seem to require that you focus on net gains and losses in good states of affairs in deciding on how you should act. But, in the case of the drowning child, we saw that part of what made your omission objectionable had to do with the your failure to respond to *that* child's imminent peril. It is not that classical utilitarianism fails to yield the right

judgment (i.e. that refusing to rescue the child is seriously morally wrong); rather, it does not adequately account for the content of that judgment. The reason is that classical utilitarianism directs our focus in the wrong direction. It directs us to see the suffering of particular individuals solely in terms of how that suffering contributes to the general state of the world. As a result, the drowning death of *that* child is significant, if it is, only by reference to how that death contributes to net gains and losses in good states of affairs.³⁴ I have argued that this is inadequate as an account of the wrongness of refusing to rescue that child and similar cases of easy rescue.³⁵ But the distinction between source and focus opens up a line of response for the classical utilitarian.

She may grant that we attach independent moral importance to the recognition of particular individuals' suffering, but argue that such psychological dispositions as these actually serve the overall good. For if everyone were to cultivate and act on dispositions to focus on the suffering of particular individuals as independently disvaluable, the likely result will be a net increase in good states of affairs when compared to alternative dispositions. Arguably, Mill himself offered this as a response to the concern that classical utilitarianism might demand unrestrained beneficence.³⁶ While the classical utilitarian is not about to give up the claim that the source of the rightness and wrongness lies in the net gains and losses in good states of affairs, she might concede that our focus, in any given set of circumstances, ought not

³⁴ These criticisms echo those made in Stocker (1976). Unlike Stocker, however, I do not despair in finding an acceptable form of (so-called) duty ethics. His insightful observations ought to function, instead, a constraint on the shape of an acceptable normative theory.

³⁵ I have not argued that this account may not be sufficient in other cases, though for reasons that will emerge in later chapters, I believe that reference to promoting the good will fail as an adequate account even in cases the classical utilitarian takes to be clear demonstrations of the view. See Chapter 5.

³⁶ See Mill (1863/2001), p. 19. See also Hare (1981), Chapter 11.

extend beyond the plight of the individuals directly affected. In this way, the classical utilitarian can account for the sense that our response to individuals ought not be mediated by the detached perspective of the impartial spectator, without giving up the central classical utilitarian idea of utility-maximization.

Unfortunately, this sort of tactic—sometimes called “indirection”—is not only internally unstable, but insufficient to deal with the original objection. While this tactic may deliver judgments that seem to cohere with our considered judgments in many of our original examples, it faces a dilemma when it comes to the cases like the expensive boots. The classical utilitarian’s desire to “have it both ways” is frustrated in this case.

Either your focus should or should not extend beyond the plight of the drowning child; either you should calculate the expected net losses in good states of affairs vis-à-vis your expensive boots or you shouldn’t. If you should, then you should allow the child to drown in hopes of saving many children abroad. But most agree that this response does not cohere with our considered moral judgments—something has gone wrong if our theory enjoins us not to save his life. But if your focus should not extend beyond the plight of the child, then while you see a reason to save the child, the question arises, what reason do you have, on the classical utilitarian account, to bring about a net *loss* in good states of affairs? After all, the classical utilitarian will continue to maintain that right action is that which produces the highest net increase in good states of affairs over bad. If the classical utilitarian is to do more than pay lip service to the concern expressed in the expensive boots case, then she is committed to the claim that you have a reason to rescue the drowning

child. But this commits the classical utilitarian to the claim that you have a reason to bring about a net loss in good states of affairs. (Incidentally, it seems implausible to claim that, in the long run, saving that child's life will recoup the losses brought about by that rescue.) So it would appear that the classical utilitarian can address the concern only by rendering her view inconsistent.

Even if the classical utilitarian could see her way around this inconsistency, there is an additional problem facing the indirection response, viz., it fails to deal directly with the objection. Recall that the objection is rooted in the concern that the wrongness of not rescuing the imperiled individual (at little cost to yourself) is not exhausted by the claim that you failed to promote the good. It is not a response to this objection to say that the additional content is supplied by the fact that you didn't promote the good indirectly.

The problem with the indirection response is that it aims to give the *impression* that there is content to the claim that your failing to assist the drowning child represents a wrong to that person, independently of net gains and losses in good states of affairs. But of course, according to classical utilitarianism, there is no content to that claim—or at least none that she recognizes. So the classical utilitarian must buttress her response with the claim that the appearance that there is more to the wrongness of refusing to aid others than promoting the good is just that, an appearance. But there is no reason, at this point, to concede this: for all we know now, there may be morally significant content to that claim. Indeed, vintage boots gives us one step in that direction. At any rate, the classical utilitarian is in the position of having to deflect each viable account of what this additional content might

be like. So, in later chapters, we will have to confront the classical utilitarian once again to see whether her doubts can be answered.

In conclusion, there seems little doubt that we have a reason—indeed, a conclusive reason—to rescue the drowning child if it is in our power, but I have tried to show that this reason seems to involve more than a reason to promote the good. Exactly what this is a subject for later chapters, where I lay out my own position with regard to reasons for rescue. For now, it should be clear that the classical utilitarian interpretation falls short of what is needed in the context of our original examples.

VI. Consequentialism

But even if classical utilitarianism cannot provide the needed analysis, perhaps a more general consequentialist approach can. After all, I have not tried to make the case that considerations having to do with consequences have no bearing on our deliberations in this context. Such a case would be difficult to make. For surely part of what makes our reason to rescue the child the conclusive reason that it is has to do with the consequences one hopes to bring about by such actions. I have in mind not only the loss of life to the child but the subsequent anguish to the parents and community. When I reason that my actions could easily prevent such things, my reason to act becomes all the more compelling. Perhaps consequentialism can provide the needed analysis by exploiting this fact.

Consider, then, the following analysis of our standing reason to render aid: we have a reason to relieve serious suffering because suffering is intrinsically bad—full stop. In Jamie Mayerfeld's (1999) words, suffering "cries out for its own abolition or

cancellation;” not only is it bad for the person that experiences it, “but bad from an impersonal point of view”(111). And it is because suffering is intrinsically bad that it ought not to occur. From this, Mayerfeld concludes that we have a standing reason (he calls it a “a prima facie duty”) to prevent it. So on Mayerfeld’s account, the independent recognition that suffering is bad leads not only to the idea that it ought not to occur, but to the idea that we have a reason to prevent it.

How might this be seen as an advance over classical utilitarianism? First, the classical utilitarian characterization of our standing reason to render aid in terms of promoting the good did not, by itself, suggest any important moral distinction between promoting someone’s happiness and preventing someone’s suffering. The classical utilitarian simply treats these as two sides of the same coin. Against this idea, Mayerfeld (1998: 130) argues persuasively that there is an important moral asymmetry between happiness and suffering, such that it is sometimes better “to bring about a smaller reduction of suffering than a larger increase in happiness, and a smaller reduction of more intense suffering than a larger reduction of less intense suffering.” The reason is that the badness of suffering is of greater moral significance than the goodness of happiness.³⁷

Second, this analysis goes some way towards meeting the focus objection discussed above. The badness of *that* child’s suffering provides you with a reason to act that is not mediated by calculations of net gains and losses in good states of affairs. It is the fact that that child’s life is imperiled, together with reasonable assumptions about what his death would mean for others, that gives you a reason to

³⁷ See also Honderich (2003). Like Mayerfeld, Honderich sees the badness of suffering as raising a special category of reasons for rescue, unlike the more generic category of “promoting the good.”

rescue him. Similar considerations hold for the other cases of easy rescue: the injured bicyclist, Kitty Genovese, and so on.

But while there is surely something to this idea, Mayerfeld, like many consequentialists, fails to illuminate how suffering gives rise to such a reason, and it is on this point that much of the discussion hinges. Remember that the second step of our schema involves identifying the standing reason that best explains our responses to our original cases. We want to know what it is about how our relationship to those individuals and/or the features under which their circumstances are described that generates a reason to assist them under the circumstances. But instead of trying to demystify the relationship between suffering and our reason to render aid, Mayerfeld adds to the mystery. On some occasions, we are told that our standing reason “arises from” or “emerges from” or “springs from” the badness of suffering. On other occasions, we are told that our standing reason is “based on” or “falls out of” or is “tied to” the badness of suffering.³⁸ But Mayerfeld’s analysis goes no further than this.

But such a crucial relationship cannot be left to metaphor. Notwithstanding the lively nature of this account (what with reasons springing from and falling out of the badness of suffering), the reader remains more perplexed than ever about what this relationship is supposed to be. The consequentialist is correct to point out that there is *something* about suffering itself that gives those of us in a position to help a reason to prevent the suffering. However, it would be a mistake to conclude straightaway that our standing reason to render aid is consequentialist in nature. For until we can illuminate the exact relationship between suffering and our reason to

³⁸ Mayerfeld (1998): Chapter 5.

prevent it, we cannot assume that bringing about better consequences represents the appropriate justification for our actions.

To his credit, Mayerfeld is forthright about his reasons for adopting a consequentialist ethics in this context. His remarks are telling.

What gives consequentialism its force is the quite specific horror of a number of specific evils, and the overwhelming need to avert them. One such evil is suffering. The vast amount of suffering in this world contributes greatly to the plausibility of consequentialism, or something close to it. If there were no suffering, consequentialism would be much less plausible (118).

Mayerfeld exposes several assumptions in this passage, assumptions that a number of consequentialists appear to share. First, he assumes that the “vast” amount of suffering in the world must be averted. But this is simply too vague to be assessed seriously. Even if we suppose that the world would be a better place if there was less suffering, on what grounds does it follow from this that you or I have an obligation to avert it? After all, the world would be a better place if tornadoes didn’t exist, but none of us has an obligation to do something to avert them. More to the point, it might be said that the world would be a better place if any number of things happened—if democratic reform swept the Middle East, if developing nations were allowed access to US markets, if farm subsidies for cotton farmers in the US were eliminated, if AIDS drugs were distributed across Africa, and so on and so on.³⁹ But none of these facts, together or individually, implies that you or I have some special

³⁹ Again, this is not to concede that statements like this one—i.e. statements that purport to evaluate “the world” or “an overall state of affairs” as good or bad—are literally true or false. The point made in the text merely assumes their truth-aptness for the sake of argument.

responsibility to bring them about. What is missing is a claim about how these facts are connected to our reason, if there is such a reason, to act.

A fortiori, these facts do not of themselves favor one moral theory over another. After all, to believe that the vastness of world suffering “contributes greatly to the plausibility of consequentialism” is already to assume that we have a reason to bring about better and better consequences—which is just what consequentialism asserts. But this is patently circular. Again, without some indication of what the relationship is between suffering (be it vast or contained) and our reason to do something about it, we can lament the state of the world while at the same time reasonably wonder what, if anything, we ought to be doing about it. There is a tendency among consequentialists to think that no argument is needed in order to move from the “vastness” of world suffering to the claim that we ought to do something to avert it. But surely non-consequentialists are entitled to such an argument. I happen to believe that the vastness of world suffering is *part* of our reason to avert it, but I do not believe that that reason is one that can simply be “read off” the nature of suffering. The notion of suffering crying out “for its own abolition,” while poetically resonant, does nothing to illuminate our reasons to prevent suffering.

In sum, classical utilitarianism and consequentialism more broadly fail to explain our reason to render aid in cases of easy rescue. As such, the grounds for the revisionist hypothesis cannot be established. Since our standing reason to render aid, if there is one, involves considerations that go beyond promoting the good, the

revisionist cannot establish what the boundaries of our moral requirements to rescue others are.

But there is another move in the revisionist schema we must consider, viz., the extent to which distant suffering counts as an instance of easy rescue. If it does, then the revisionist might legitimately insist that we have a standing reason to render aid to distant sufferers, even if the standing reason involves considerations that go beyond issues of promoting the good. Independent of the content of this standing reason, however, I will argue in the next chapter that distant suffering does not meet the criteria set out in *ER*. And this makes revisionism all the more difficult to maintain.

CHAPTER THREE

Easy Rescue and Distant Suffering: A Closer Look at Revisionism

MUCH OF THE RENEWED INTEREST in questions of positive duties has followed the work of people like Peter Singer and Peter Unger, whose emphasis on distant suffering has in many ways changed the tenor of the debate. To be sure, there is nothing new about extreme poverty (aside perhaps from greater awareness of both of the extent and degree of such poverty). Nor is the question of beneficence new. After all, we trace the parable of the Good Samaritan back to the New Testament.⁴⁰ What *has* changed is the extent to which even minimal efforts on the part of well-off individuals like you and I can actually save lives. In their classic work on famine, Jean Dreze and Amartya Sen (1989: 3-4) write:

Hunger is not a new affliction...Hunger is, however, intolerable in the modern world in a way it could not have been in the past. This is not so much because it is more intense, but because widespread hunger is so unnecessary and unwarranted in the modern world. The enormous expansion of productive power that has taken place over the last few centuries has made it, perhaps for the first time, possible to guarantee adequate food for all, and it is in this context that the persistence of chronic hunger and the recurrence of virulent famines must be seen as being morally outrageous and politically unacceptable.

Whether or not the persistence of chronic hunger is “morally outrageous,” there can be no question that our current ability to alter the conditions of starving individuals half a world away changes the moral landscape. For, apparently, I can now save (very many) lives through modest changes in my life, where this was simply not possible in the past.

⁴⁰ Luke 10:25-37.

In the age of modern humanitarian-relief agencies like UNICEF, Oxfam, and CARE, the marginal utility of your one dollar in places like Eritrea or Bangladesh is unprecedented.⁴¹ It is this empirical fact—what I will call the fact of modern humanitarian aid—that has added a new degree of urgency to the debate. No longer does the question of beneficence simply concern how much I ought or must do for the underprivileged, say, in my own community. Now the question concerns how much I ought or must do to *save people's lives*. It is in the context of the fact of modern humanitarian aid that Singer and Unger assimilate the case of distant suffering with the more dramatic forms of local suffering.

In this chapter, we consider the case for revisionism as it applies to distant suffering. In particular, we will look at the question of whether distant suffering counts as a case of Easy Rescue (*ER*). I suggest a number of reasons why it should not count as a case of easy rescue, at least according to the model that covers cases of local emergencies. This should not imply that there won't be strong moral reasons for contributing, even if they do not rise to the level of the strength of one's reason to respond to local emergencies.⁴²

I. Distant Suffering as Easy Rescue: A Defense

Consider, again, the conditions I outlined in the previous chapter, conditions that together constitute what I called easy rescue, or *ER*.

⁴¹ The details, of course, are controversial. But while we can dispute how much it costs to save an individual's life, there is broad consensus on what your contributions can purchase. All of the major humanitarian relief agencies provide detailed accounts of just where and what money is spent on: see, e.g. www.oxfam.org, www.unicef.org, and www.care.org. Widely considered to be the standard reference on the effectiveness of international aid efforts is Cassen (1994). Unger (1998) provides a summary of much of these findings.

⁴² See Chapters 6 and 7.

- f) an innocent person stands to suffer a significant loss, like a limb or his life;
- g) some agent *S* is in a position to prevent such a loss, either through her actions or her resources;
- h) *S*'s executing those actions or contributing those resources can be expected to result in only a trivial cost to *S* when compared to the loss the victim faces;
- i) *S*'s efforts at rescue under the circumstances do not require *S* to perform actions that are prohibited by commonsense morality; and
- j) *S*'s efforts at rescue are not likely to increase the number of subsequent individuals needing rescue.

Let me start, first, with reasons to think that these conditions *do* in fact cover the case of distant suffering. It would be hard to deny that (very many) innocent individuals currently face a serious loss.⁴³ As Singer (1972: 243) rightly noted more than thirty years ago, “the facts about the existence of suffering are beyond dispute.”⁴⁴ They still are. But what of condition (b)? The phrase “in a position to prevent such a loss” is of course vague—but for good reason. For we needed a way of capturing the various scenarios in which it seemed as if *S* was under an obligation to prevent suffering. If you’ll recall, “preventing suffering” took on various shapes depending on the scenario; in some cases, it meant pulling a child from a shallow pond, in others, driving someone to the hospital, in others, giving up a trunk full of clothes. In some cases, it refers simply to calling 911. And obviously we could add to this list. But so long as we remain at this level of generality, there does not seem to be anything linguistically odd about the suggestion that I am “in a position” to save a distant person from a serious loss—assuming of course the fact of modern humanitarian aid.

But can we expect *S*'s contributions to result in only *trivial* decreases in *S*'s well-being when compared to the loss the victim face, as (c) asserts? This is a more

⁴³ See Appendix.

⁴⁴ Singer, *op. cit.*: p. 243.

difficult question because, unlike the cases we have so far considered, there are very many individuals needing rescue in the case of distant suffering (on the order of millions). Hence, it might be thought that *S* stands to lose a great deal were she to address the full scope of the problem. *S* would find herself bankrupt (or worse) well before she prevented all the suffering that currently exists. But, it might be responded, this concern flows from a misreading of (c). For (c) does not make reference to the *aggregate* cost of rescuing all individuals in peril, but only the cost incurred from rescuing this *one individual* here. Perhaps aggregate cost should make a difference, but as far as (c) is concerned, it does not. (c) instructs us to compare the cost to me of saving one life to the cost to the child of suffering painfully and dying. And can we really say that parting with a few dollars will pain me as much as starving to death will pain that child in Eritrea? Unless we have in mind very exotic cases, it is hard to think it will. This leaves us then with conditions (d) and (e), which should require only a brief discussion.

True, there are cases in which contributing a few dollars would violate other commonsense moral prohibitions. For example, someone who finds a wallet and instead of returning it to its rightful owner sends the three dollars it contains to a relief agency has violated, one might argue, a commonsense moral prohibition. Likewise for someone who acquires that money through acts of manipulation, assault, exploitation, or the like.⁴⁵ But these cases raise separate questions, questions that are not immediately relevant to the present discussion. For the question we're

⁴⁵ Unger (1995) argues that one is morally justified in stealing from a millionaire and contributing the funds to Oxfam, so long as one can avoid detection (see his Chapter X). At least one can say that Unger accepts the logical consequences of his moral commitments. It is views like Unger's that give revisionism its name. Needless to say, I reject Unger's claim—and the moral grounds on which it rests.

considering is, is my contributing some of my disposable income to relief agencies, under the most “boring” conditions, in the least bit ethically problematic? After all, what could possibly be wrong about my sending some of my (legitimately acquired) money—which I would otherwise spend on non-essential goods and services—to a relief agency whose aim is to prevent premature deaths abroad? Again, the answer seems as obvious as it is boring: nothing.⁴⁶

II. Determinacy and Distant Suffering

Now I want to consider ways in which *a – e* fail to capture the case of distant suffering. Here I will argue that there are a number of ways that the case of distant suffering departs in morally significant ways from cases of local emergencies.

In order to get the case of distant suffering to fit comfortably under condition (*a*), one has to make several important assumptions. First, one has to assume that the move from a single individual to a sprawling collection of individuals does not alter one’s moral relationship to the potential recipient(s). One has to assume, in other words, that there are no notable differences in the fact that a single, determinate individual is imperiled in one case and that a loose, vaguely specified group is imperiled in the second. But surely there are important differences.

One difference that I think is morally significant concerns determinacy. In the move from a single, determinate individual to a loose, vaguely specified group, our reasons to respond—and with it, the stringency of the requirement to respond—are

⁴⁶ We might legitimately require that individuals, if they intend to contribute, make “responsible contributions,” where these are secure contributions (e.g. checks, money orders) to established organizations which are not only free from corruption themselves, but avoid perpetuating corruption abroad. If any organizations count as established, Oxfam, UNICEF, and CARE surely do.

altered. In the case of the drowning child, for example, one can easily identify the recipient of one's actions (or, inactions, for that matter): that child there. This is not at all the case when we move to the case of distant suffering. The structure of relief agencies like UNICEF and Oxfam is such as to obliterate the connection between you (or, more specifically, your contribution) and the recipient.⁴⁷ The trick is to show why this is important. Part of the reason has to do with claims made in the previous chapter: our moral reasons to respond beneficently encompass more than simply what sorts of ends we bring about; they also encompass how our moral *relations* to others are affected by our actions.

Attempts to show that responding to distant suffering counts as a case of easy rescue assume that there is a determinate (if not determinable) individual who stands to gain or lose by your actions. This was entailed by the first criterion laid out above. But this assumption is false. What is more, that there is no determinate individual makes a critical difference in our moral relationship to distant sufferers. As I argue below, the concern to prevent a determinate individual from losing her life is distinct from the concern to simply decrease the number of lives lost by one, where there is no fact of the matter as to who this one individual is. And the distinction turns on the sort of relation one stands in to potential recipients.

Singer has long advocated strict duties of humanitarian relief, arguing that most of us, on reflection, will recognize that we are likewise committed to such duties. Working backwards, I want to reconstruct Singer's argument in a way that exposes the moral importance of determinacy. First, according to Singer, a

⁴⁷ This claim invariably attracts skepticism, much of it unwarranted. First, I do not mean to doubt that when individuals such as Bill Gates contribute \$10 million dollars to AIDS relief there is no fact of the matter as to who benefits. The question is whether there is a fact of the matter as to who benefits when I contribute \$50 to UNICEF, an organization whose annual charitable expenditure exceeds \$1 billion. That indeterminacy surrounds the sorts of contributions Singer and others ask us to make is clear once one understands the structure and operating procedures of the more reputable humanitarian agencies, a task I cannot possibly take on here. See Cassen (1994); Pogge (2002); and Dreze and Sen (1989).

commitment to strict duties of humanitarian relief follows from a commitment to the general principle

(P) If we can prevent something bad without sacrificing anything of comparable significance, then we ought to do it.⁴⁸

Since contributing to humanitarian relief organizations like Oxfam and UNICEF can lessen the suffering abroad, and since the cost to us of making such contributions is not comparable to the suffering of those abroad, we ought to contribute to organization like Oxfam and UNICEF. But on what grounds are we committed to P? According to Singer (1993: 239), P is posited as “a plausible principle that would support the judgment that I ought to pull the [drowning] child” from the shallow pond. Here Singer is appealing to our intuitions regarding the following case: On your way across campus to give a lecture, you come across a child drowning in a shallow pond. Although rescuing the child will require getting muddy and missing your lecture, most people judge that you ought—indeed, must—rescue the child. As justification for this judgment, Singer cites P. But surely there are other principles that would support the judgment, principles that are equally as plausible as P, if not more so. Indeed, if our aim is to explain our judgment under the circumstances, P fails to capture at least one important feature of the case.

As far as I can tell, the following principle is just as plausible as P, if not more so, in the context of the drowning child:

⁴⁸ Singer (1993), p. 239.

(P*) If someone *S* stands to lose something significant (like her life) without your assistance, then you ought to assist *S* when doing so does not require sacrificing anything of comparable significance.⁴⁹

After all, as far as the details of the case are laid out, you are not called upon to simply “prevent something bad;” you’re called upon to save *this child’s life*, a life that will be lost without your assistance. Of course if P represented the ultimate justification for our judgment that you ought to pull the child out, then P* would itself stand in need of justification. It would not be enough to point out that you are preventing the child from losing his life; one would need to justify why we ought to prevent others from losing their lives. But such a justification is hardly necessary, at least in cases like the drowning child. The fact that preventing someone from losing his life is an instance of preventing something bad adds nothing to our moral reasoning, and there is little to motivate the claim that it is *because* preventing someone from losing his life is an instance of preventing something bad that one ought to rescue the drowning child.

Before pushing forward with this dispute, let me emphasize why it matters, for one might wonder whether there is any interesting difference between P and P* such that we should care which principle best supports the judgment. Recall that Singer’s argument for stringent duties of humanitarian relief turns on a commitment to P. Since P makes reference only to “preventing something bad,” one can draw on empirical evidence about the effectiveness of humanitarian relief efforts to argue that contributions to organizations providing such relief in fact prevent something bad.

⁴⁹ Note the similarity between P* and the first criterion in Easy Rescue, labeled *a*.

By contrast, P* makes reference to a determinate individual (*S*) and how one's actions bear on *S*'s fate; there is, as it were, a fact of the matter as to who stands to gain (or lose) by one's response. Now here's why the contrast matters: a commitment to P*, unlike a commitment to P, does not yield stringent duties of humanitarian relief since, given contingent facts about the structure of humanitarian relief organizations like Oxfam and UNICEF, there is no determinate individual who stands to gain or lose through one's contributions.

At least in the realm of modest contributions to any of the major relief organizations, there is no fact of the matter as to who, if anyone, I save. What this means is that, for any determinate individual (*D*) in need of humanitarian assistance, *D* does not live or die according to my decision to give. Even supposing that I contributed \$100 to Oxfam last month, it would be false to say that *D* would be dead today had I refused to give, just as it would be false to say that *D* would be alive today had I contributed. The reason should be obvious: my contribution enters a general fund, from which charitable expenditures, not to mention administrative costs, are drawn. So, for example, Oxfam, whose annual charitable expenditure exceeds \$1 billion, operates in a way that renders unanswerable questions like, but *which individuals* were saved by my \$100 contribution? The claim is not that a \$100 contribution makes no difference; it is that the difference cannot intelligibly be measured in determinate lives saved.⁵⁰

⁵⁰ What we are doing when we contribute \$100 to Oxfam is increasing Oxfam's ability to carry out its mission, to address whatever need they deem critical and to address it in a way that they deem is effective. This is not unlike the way we think of our tax dollars: someone who failed to pay her taxes would have difficulty identifying which government program was under-funded as a result of her refusal to pay her taxes; the reason is, there is no such determinate program.

Returning once again to Singer's project, we can spell out the dilemma this way: given Singer's ultimate aim of defending stringent duties of aid, he has reason to defend the more amorphous P, for P appears to support the kinds of acts that can only be described amorphously, like contributing to humanitarian relief efforts. But P* doesn't support such acts because P* makes reference to determinate beneficiaries. And yet, as I've argued, P* is a more plausible principle when it comes to why I ought to pull the child out. So while P would support Singer's ultimate conclusion, the case of the drowning child does not uniquely favor P. And so Singer needs another case in support of P, one that would uniquely favor P.

I believe that I can provide Singer with such a case. In the next section, I describe a case that appears to be supported by P and thus appears to support Singer's aim of establishing stringent duties of humanitarian relief. What this case does not establish, however, is Singer's (1993: 230) claim that "we have an obligation to help those in absolute poverty that is *no less strong* than our obligation to rescue a drowning child from a pond." This claim would only follow if determinacy did not make a moral difference. But I offer a second, amended case that strongly suggests that determinacy does make a moral difference.

Detention Center I. One way of supporting a principle like P—one that asserts that if we can prevent something bad without sacrificing anything of comparable significance, then we ought to do it—would be to consider a case like the following.⁵¹

⁵¹ I am grateful to Peter Carruthers who first suggested to me the basic outline of this example.

Terrorists are about to strike a youth detention center where 2,000 youths are detained, each in a separate cell. As it happens, cell doors can only be opened electronically, from a centralized computer. As part of the attack, however, the terrorists have disabled the computer in such a way that only 999 cell doors will be opened automatically before the center's destruction. For unknown reasons, however, the computer will select the 999 doors *at random*.⁵² So until the designated moment at which the randomizer selects which doors to open, there is no fact of the matter as to which doors will open and, hence, which youths will be saved. There is, however, something you can do in the little time remaining: you can enter into the computer a lengthy code that will instruct the system to open 1,000 cell doors, instead of 999—though, again, which doors will open will be the result of the randomizer.

This case is unusual in several respects. First, although entering the code will result in one fewer deaths than not entering the code, there is no determinate individual whom we can point to as the individual whose death you prevented. After all, for any one youth whose door was opened in the actual world, there are many near possible worlds in which your entering the code did not result in his door being opened. Thus no individual can claim that his death was prevented as a result of your decision to enter the code, since, for all that anyone knows, his cell door would have opened even without your assistance. Correlatively, no individual can claim that his life was lost as a result of your decision *not* to enter the code, since for all anyone knows, his cell door wouldn't have opened even with your assistance.

⁵² To ensure metaphysical indeterminacy, imagine that the randomizer involves, say, emissions from a small radium source.

Another curious feature of this case is that despite the obvious significance of your decision to enter the code (viz., preventing one fewer deaths), the most that any one individual can claim is that by entering the code you increase his chances of being saved by only a fraction. For any youth, the probability of his being saved without your assistance is 0.4995; with your assistance the probability increases to 0.5, a difference of a mere 0.0005.

But surely no one would want to conclude on the basis of these considerations that you have no reason to enter the code. Rather, these considerations suggest that your reason to enter the code is essentially *aggregative* in nature; ceteris paribus 1,000 deaths is “less bad” than 1,001. But unlike some other forms of moral dilemma that invite aggregation, **Detention Center I** is distinctive in that either way you choose to respond, there is no determinate beneficiary (if you choose to enter the code) or determinate victim (if you refuse to enter the code)—even *after* your decision has been made. What we know is that either 1,000 individuals evacuate the center or 999 do. What we do not know, nor could know, is who that additional youth is since there is no fact of the matter as to the identity of that individual. Apparently, then, it is only at the level of the aggregate that we can appreciate the strength of your reason to respond.⁵³

Moreover, we now have a case that appears to support P. For most people judge that you ought to enter the code, and a plausible way of accounting for this judgment is P: if you can prevent something bad without sacrificing anything of

⁵³ As I discuss later, however, it is perfectly legitimate to point out that the effort required to enter the code is so minimal as to make omission unreasonable. The point, however, concerns the asymmetry between the strength of your reasons to respond when considered from the impersonal perspective, wherein aggregative facts are salient, and from the perspective of any one (potential) beneficiary: seen from the impersonal perspective, your reason appears far more stringent.

comparable significance, then you ought to do it. That there is no determinate victim (or beneficiary) in **Detention Center I** is precisely why we speak of “preventing something bad” instead of “saving someone’s life.” As far as P is concerned, it is enough that a) 1,000 deaths is “less bad” than 1,001 deaths and b) entering the code does not require any comparable sacrifices.

Detention Center II. Now all of this is precisely what we would expect if determinacy possessed no independent moral weight. But now consider the following alteration to our example.

Next to the computer that controls the 2,000 cell doors is a second computer that controls a single cell door, separate from the rest, and which is unaffected by the operations of the main computer. There is no indeterminacy about who will be evacuated if you decide to enter a different (though equally lengthy) code into this second computer, viz., the occupant of cell #2000, say.⁵⁴ Assume as before that even without your assistance 999 cell doors in the main detention center will open automatically. Now imagine that you only have enough time to enter one code into one computer. What should you do?

One might immediately doubt whether much hangs on your decision. After all, whether you enter a code into the main computer (Option 1) or enter a separate code into the second computer (Option 2), the outcome is the very same: 1,000 lives saved (and 1,000 lives lost). Option 1 ensures the survival of 1,000 youths (selected

⁵⁴ To keep the numbers straight, imagine that the main computer now controls 1,999 cell doors, instead of 2,000. In both examples, then, there are exactly 2,000 individuals whose lives are at stake. Of course this change is not entirely without consequence: now, for any youth in the main detention area, the probability of his being saved without your assistance is now 0.4998; with your assistance, 0.5002.

randomly) and the death of the occupant of cell #2000 and 999 other randomly selected individuals. Option 2 ensures the survival of #2000 and 999 randomly selected individuals and the death of 1,000 randomly selected individuals. The only difference between the two options is one of determinacy: on Option 1, there is no determinate beneficiary, even though 1,000 (as opposed to 999) randomly selected individuals are rescued. On Option 2, there *is* a determinate beneficiary, even though only 999 (as opposed to 1,000) randomly selected individuals are rescued. Negatively speaking, on Option 1, there is a determinate *victim*—viz., #2000; on Option 2, there is no determinate victim.

Now if determinacy did not make a moral difference, if, that is, the determinacy of a potential beneficiary (or victim) should not be taken into account in deliberating what you ought (morally) to do, then you should be indifferent as to which option you take since, apart from determinacy, the options are effectively equivalent. But I submit that you should *not* be indifferent as to which option you ought to take: by choosing Option 1, you bestow a miniscule benefit onto any one occupant of the main detention center *at the cost of #2000's life*, and this amounts to a refusal to recognize the special claim #2000 has on you. And this, I argue in later chapters, amounts to a moral failure.

Let us first concede that every individual in the detention center has a reason to want to avoid serious harm; this much is agreed. The question concerns whether there are considerations that raise any individual's or any group's reason for rescue above others' reasons, in the sense that he or it has a greater moral claim to your assistance. If there are such considerations, then we should not be indifferent as to

which option you take. Three possibilities present themselves. First, one might assert that any individual in the main detention center has a greater claim to your assistance than the occupant of cell #2000. I see no plausible case to be made for this possibility. For if the view is that any individual in the main detention center has a greater claim on your assistance because, in general, one has a moral reason to want to increase his chances of survival, then the occupant of cell #2000, if anyone, should have a greater claim on your assistance since Option 2 increases his chances to 1.

A more plausible possibility is that the *group* has a greater moral claim to your assistance than the occupant of cell #2000. Together, they can assert that Option 1 will guarantee the survival of one more individual. But this view supports the hypothesis that one should be indifferent as to which option you take, not the hypothesis that you should place greater moral weight on one entity over another. After all, the occupant of cell #2000 has the identical claim: Option 2 will guarantee the survival of one more individual. The only difference, of course, lies in the identity of the surviving individual.

One might attempt to bolster support for the view that you should be indifferent as to which option you take by claiming that in choosing Option 2 you treat the group unfairly, in this sense: you are putting one individual's claim for assistance above the claims of 1,999. Since these individuals, taken together, appear to have the same claim as the occupant cell #2000—viz., that the option favoring them will result in one less death—it would be morally arbitrary to choose the individual over the group. After all, what sense does it make to attribute special moral weight to the claims of the individual when the group has the same claim

against you? And vice versa. As tempting as this line of argument is, it faces serious obstacles.

The first obstacle, which I discuss at length in Chapter 5, concerns the suggestion that individuals can “aggregate” their individual complaints in such a way as to yield a *net complaint*, that is both distinct from and, in this instance, “weightier” than their individual complaints. This is a familiar line, and it is undermined by familiar objections. It would justify, for example, certain limited forms of slavery; it would justify bringing severe harm to an individual (including death) for the sake of relieving many, many minor inconveniences, such as the headaches of a million people. The problem, which I won’t spell out here, lies in a conception of value according to which the ultimate source of value consists of states of affairs considered impersonally. This conception grounds the possibility of aggregating the value or disvalue of individual well-beings.

If I can persuade you that such a conception of value is mistaken, and here I can only issue a promissory note, then what we are left with is the question of whether or not any *individual* has a greater claim to your assistance. What is ruled out is the possibility that the group could coherently have a claim (let alone, a special claim) to your assistance. I have also ruled out the suggestion that any individual in the detention center has a special claim on your assistance. This leaves two possibilities: either no one has a special claim or the occupant of cell #2000 has such a claim. It is worth asking: first, what sort of complaint does any individual in the main detention center have to your choosing Option 2? Second, what sort of complaint does #2000 have to your choosing Option 1?

As we've already conceded, every member of the detention center has a reason to want to avoid death when it is in your power to do so; therefore, in that sense each has a claim on your assistance. By choosing Option 2 you will not be honoring any one of these individual's claims. But bear in mind what this claim is a claim to—viz., to increase his chances of survival by 0.005. *Ex hypothesi*, this is the *most* you can do for him; there are substantial, external constraints on the shape your assistance can take. It is thus false (and dangerously misleading) to suppose that his claim is a claim to anything more. Still, it is true that by choosing Option 2 you decrease his chances of survival.

Like the other members of the detention center, #2000 has a reason to want to avoid death and so in that sense he, too, has a claim on your assistance. By choosing Option 1 you will not be honoring his claim. But #2000's predicament is altogether different from the predicament of any other member, for you have the option of saving his life, not merely increasing his chances of survival. By choosing Option 1 you allow him to die. It is plain then what #2000 is demanding when he demands that you choose Option 2: It is a demand that each person allows his chances of survival to decrease by 0.0004 *so that* one person may live. We can grant, therefore, that #2000 has reasonable grounds for demanding that you choose Option 2, *even if* those grounds are not in the end sufficient.

Now what would any one in the main detention be demanding in demanding that you choose Option 1? He would be demanding that you allow #2000 to die *so that* his chances of survival could be increased by 0.0004. It is here that the case becomes clearest. This is not a reasonable demand. I decrease my chances of

survival every time I get into an automobile or fly in an airplane or, for that matter, enter a smoky bar or indulge in artery-clogging foods. I often accept such risks (as I assume my reader does, too) for minimal rewards. But if decreasing our chances of survival by a fraction meant that much to us, if we were indeed prepared to allow another person to die in order to avoid such risks, it would be altogether mysterious why we would so frequently and so willingly engage in activities that had the effect of decreasing our chances of a long life when doing so was quite unnecessary. But it is not mysterious at all. There are innumerable reasons why we accept such trade-offs; life is too important—too rich in experience and experiment—to cut oneself off from any risk. But if our reasons for accepting a decrease in our chances of survival can include such things as aesthetic and intellectual considerations (to name just two), then surely they could include someone's losing his life as a relevant consideration. This is not the question of making a career out of saving lives; it is a question about this token case, about one's deciding between Option 1 and 2.

The problem with indifference in this case is that you are assigning equal weight to the claims of every member in the detention center, including #2000, but there is a critical difference between demanding that someone decrease his chances of survival by a fraction of one percent so that someone else may live and demanding that someone be allowed to die so that another can increase his chances by a fraction of one percent. The difference lies in your relationship to the individual who has the most to lose (and gain), and this points to a theme to be developed in Chapter 6 and 7.

That theme concerns a more general feature of morality that the present case touches on, viz., an ideal of human relationship. Very crudely put, the morality of

right and wrong does not characterize one's special (maximizing) relationship to the world; the morality of right and wrong characterizes one's relationship to other persons. Relationships of this latter sort represent, in the words of Michael Thompson, a "bipolar" or reciprocal normativity.⁵⁵ What this means is that your reasons to respond to #2000 are directly tied up with *his claims against you* and the fact that he is specially vulnerable to being wronged by you under the circumstances. What is more, #2000 has a claim on you that none of the other youths has on you: your actions *directly determine* his fate.⁵⁶ For him, the difference between Option 2 and Option 1 is the difference between life and death. For any other detainee, the difference between Option 2 and Option 1 is a tiny fraction of 1%.

In contrast to **Detention Center I**, where your reason to respond appeared to derive from considerations of aggregation (1,000 lives saved versus 999), **Detention Center II** introduces a reason to respond that is independent of such considerations. It is a reason that involves reference to how your response will affect the fate of a determinate individual. And this, it seems to me, fundamentally alters your moral relationship to those who look to you for assistance.

Picking up the thread of our earlier discussion, we can now see that justifying Option 2 requires more than simply an appeal to "preventing something bad," for both Option 1 and Option 2 will prevent something bad. A justification for Option 2 requires appealing to how one's refusal to respond affects a determinate individual's fate (in this case, #2000). One might then appeal to P* in justifying one's decision;

⁵⁵ See Michael Thompson, "What is it to Wrong Someone? A Puzzle about Justice," and Phillip Pettit and Michael Smith, "The Truth in Deontology," both in Wallace, Pettit, Scheffler, and Smith, eds., Reason and Value: Themes from the Moral Philosophy of Joseph Raz (Oxford: Clarendon Press, 2004). See also Wallace (forthcoming).

⁵⁶ This idea is spelled out fully in Chapter 6.

one might claim, in other words, that if someone *S* stands to lose something significant (like her life) without your assistance, then you ought to assist *S* when doing so does not require sacrificing anything of comparable significance. **Detention Center II** then forces us to decide between P and P*; it forces us, in other words, to weigh the moral significance of determinacy since this is effectively what distinguishes the two principles.

What larger implications can we draw from this? First, however one chooses to describe the wrong one commits by refusing to contribute to humanitarian relief, that description will lack a crucial element, one which is involved in the description of the wrong one commits by refusing to rescue the drowning child—viz., that a determinate individual lost his life as a direct result of your omission. Contingent facts about humanitarian relief organizations render one's beneficiaries indeterminate. It is a related mistake to ground our obligations to contribute to humanitarian relief on our obligations to respond to local emergencies. For the drowning child has a claim on you that no distant sufferer does. By refusing to rescue the child, you wrong him in a way that you cannot wrong a distant sufferer.⁵⁷ This is not the claim that you cannot wrong distant sufferers. Rather, it is the claim that the sort of wrong you commit by not contributing to humanitarian relief, if it is a wrong, can only be understood by reference to “preventing something bad” where this is not decomposable into the claim that this or that individual lost his life as a result of your refusal to give. It would be in vain then (not to mention, beside the point) to search for the individual who perished as a result of your refusal to contribute last month.

⁵⁷ I explore this difference in greater detail in my “Good Samaritans, Good Humanitarians: The Case for Unique Dependence,” under review.

There is no such individual. But Singer-style arguments derive much of their force from the assumption that there is.

To summarize, the potential beneficiaries of humanitarian relief do not stand in the same sort of relationship to us as the drowning child does, for our relationship to the child is, *ex hypothesi*, a relationship to a determinate individual. Admittedly, I have done nothing yet to illuminate the moral significance of this relationship, but it is nonetheless a feature we are willing to draw into our moral deliberations. At any rate, my argument should expose the mistake in Singer's claim that "we have an obligation to help those in absolute poverty that is *no less strong* than our obligation to rescue a drowning child from a pond." This claim would hold only if determinacy did not make a moral difference, but, apparently, it does make a difference.

These results should not imply—and I can't emphasize this enough—that you have no reason to contribute to humanitarian relief. You have a reason, just as you have a reason to enter the code into the main computer in **Detention Center I**: you increase the number of lives saved; you bring about a better state of affairs. But this claim does not license the claim that there is a determinate individual who benefits from your efforts; this is the move I've been attempting to block. So while it may be legitimate for me to criticize you for your failure to contribute to humanitarian relief (or so one might think), it would be *illegitimate* for me to justify my criticism on the grounds that this or that person perished as a result of your inaction. That is a different sort of criticism. And a more serious criticism, for it would be wrong of me to be indifferent as to what I should do when faced with the choice of aiding a

determinate beneficiary or aiding an indeterminate beneficiary.⁵⁸ This is the result we saw in comparing Option 1 and Option 2: there is a more stringent reason to enter the code into the second computer, where this will have the effect of opening cell #2000. And this intuition is quite robust.

If **Detention Center I** is an appropriate analog to the case of humanitarian relief, at least as far as determinacy is concerned, then just as my refusal to enter the code reduces (for any determinate detainee D) D's chances of survival by the merest of fractions, so, too, my refusal to contribute \$100 reduces (for any potential recipient of aid R) R's chances of survival by the merest of fractions. And just as there are many near possible worlds in which D's cell door does not open up *even though* I entered the code, so, too, there are many near possible worlds in which R does not receive life-saving aid *even though* I contributed \$50. Before going on to connect the issue of determinacy to a larger model of moral reasoning, I want to address two potential concerns that determinacy seems to raise.

DETERMINACY AND GUILT. It might be asked whether #2000 is aware of your choices as described in **Detention Center II**, for if he is unaware of your choices, then it may not strike one as intuitively obvious that you have a more stringent reason to choose Option 2 over Option 1. The thought here is that our intuition in **Detention Center II** does not track determinacy per se; instead, it tracks the emotional difficulty of refusing to rescue #2000 when he knows the options available to you, that is, when he knows that *you* decide *his* fate. Lacking that knowledge, however, he cannot look to you for special consideration, nor express gratitude (or resentment) upon your

⁵⁸ See Gomberg (2002).

making a decision. Thus one may think that the differences between Option 1 and Option 2 are merely differences in our sensibilities, and you could (and, perhaps, should) be indifferent as to which option you choose since the outcomes are, morally speaking, the same. At any rate, if this is the correct diagnosis of our intuitions, then determinacy does not appear to do any of the moral work I have suggested.

This worry is inconclusive. First, simply as a matter of psychology, it's far from clear that your emotional reaction to choosing Option 1 over Option 2 (thus allowing #2000 to die) would turn solely on what he knows or doesn't know. After all, remorse or regret can be just as acute when those we have injured or allowed to be injured are unaware of our complicity. The reason is that, apart from what others know or don't know, our attitudes toward ourselves also depend on what *we believe we've done* or are about to do.⁵⁹ It is thus relevant how we describe to ourselves our conduct under the circumstances, and this will consist in part in describing our reasons for choosing this course of action over that course.

Second, even if the differences between Option 1 and Option 2 appear to lie only in our emotional responses, this does not mean that no inference can be drawn about the nature of the situation to which we are responding. Indeed, it would be puzzling if this were not so. Why, that is, would we experience a negative emotional reaction to choosing Option 1, where #2000 is cognizant of our choices, if we did not believe that our conduct warrants criticism? If we believe that the only difference between the two options lies in our emotional responses, then we would not look to

⁵⁹ We need not imagine ourselves under the watchful gaze of a vengeful God, nor do we need to hold to an especially rigorous personal ideal. It is enough that we hold a minimal standard of moral decency, according to which we believe, by and large, that our conduct meets the standards that others could accept.

the facts of the situation as a way of explaining why we are having the reactions we are having. But surely we do look to the facts; at the very least, we seem sensitive to how #2000 would react to our decision. Indeed, I think a plausible case could be made for the claim that it is precisely because he has legitimate grounds for objecting to Option 1 that we would feel the way we would feel upon choosing Option 1. But if the objection we've been considering is correct, then such a feeling is irrational. In any event, this objection cannot be made to work unless it can provide a case for thinking that the feeling is *in fact* unwarranted, but to do this requires producing a separate normative argument showing why determinacy cannot make the moral relevance I am claiming for it. This objection does not begin to do this. Perhaps the following argument can.

VICTIMIZATION WITHOUT VICTIMS? If determinacy makes a moral difference, such that we have more stringent obligations to determinate individuals than groups, then we should expect that determinacy should also make a difference when it comes to producing harms. But there seems to be no moral difference between the wrongs committed in the following two cases.

Detention Center III. Same as **Detention Center II** except that: Sadist X enters a code into the second computer that has the effect of *locking* cell door #2000. Now, instead of being released with the other 999 individuals (as originally programmed), the occupant of cell #2000 is killed in the ensuing attack.

Detention Center IV. Same as **Detention Center II** except that: Sadist Y enters a code into the first computer that has the effect of locking 1,001 randomly selected cell doors instead of only 1,000 (as

originally programmed). Now, instead of 1,000 individuals being released (the occupant of cell #2000 is scheduled for release), only 999 individuals are.

Since Sadist Y's action seems no less wrong than Sadist X's, even though there is no determinate victim of Sadist Y's action, it seems to follow that determinacy does not make the moral difference I've been claiming.

This objection trades on the assumption that if determinacy matters in the realm of preventing harms, it should also matter in the realm of producing harms. If reasons for preventing harm can be affected by the determinacy of one's beneficiaries, then one's reasons against producing harm should, by parity of reasoning, be affected by the determinacy of one's victims. Thus we should find Sadist Y's action *less wrong* than Sadist X's action. But we don't, so determinacy must not affect our reasons in the ways I've suggested. This conclusion, however, follows only if we accept the conditionals above. If there are reasons for denying the claim that preventing harms and producing harms constitute reasons of the same moral stringency, then there is no reason why we cannot accept the judgment that Sadist Y's action is no less wrong than Sadist X's while, at the same time, claiming that morality requires that you choose Option 2 over Option 1 in **Detention Center II**. This would be one way to mount a response to the objection.⁶⁰

But, perhaps, such a defense is unnecessary. Perhaps we were too quick in accepting the intuitions that allegedly support the objection: are we sure that

⁶⁰ Presumably, strict consequentialists do not see a principled difference between preventing harms and producing harms, whereas non-consequentialists do. So my dispute is, first and foremost, with consequentialists. Those who are willing to allow that there are considerations that may figure in one domain (e.g., preventing harms) but not in the other may wish to skip ahead to section 5.

determinacy does not affect our judgments in **Detention Centers III** and **IV**? I will argue that it does; indeed, I will argue that our judgments in fact mirror our judgments in **Detention Center II**. It should be born in mind, however, that my response concedes to my critic what I in fact reject, viz., that reasons for preventing harm are no more or less stringent than our reasons against preventing harm. But I believe one can respond to this criticism without mounting a full-blown refutation of that idea. So how should we understand our intuitions to **Detention Centers III** and **IV**?

First, our tendency to see no moral difference between the actions of Sadist X and Sadist Y can be attributed to our judgments regarding their *motives*: they seek to harm others. (They're sadists, after all.) That they employ different methods seems irrelevant against the backdrop of their deplorable motives, and we are not apt to draw distinctions where such distinctions are overshadowed by undisguised maliciousness.⁶¹ One may be tempted to say the same thing with regard to **Detention Center II**: the differences between Option 1 and Option 2 seem irrelevant against the backdrop of their beneficent motives. But when we are asked to *choose* between Option 1 and Option 2 (as we are in **Detention Center II**), we are prepared to look more closely at the details of the cases and to draw into consideration facts that may not have seemed salient on a first pass, where seeking to save as many lives as possible seemed the only relevant consideration. Perhaps, then, **Detention Center III** and **IV** ought to be re-described as different options from which one is required to choose. Imagine, then, that Sadist Z offers you the following options:

Option 3. Sadist Z kills *all* the members of Detention Center and you.

Option 4. You perform the actions of Sadist X, locking cell door #2000.

⁶¹ See Rachels (1992).

Option 5. You perform the actions of Sadist Y, instructing the computer to lock 1,001 (randomly selected) doors instead of 1,000, as originally programmed.

Clearly, you have most reason to avoid Option 3. Should you, however, be indifferent as to the other options? Do you have any more reason to choose Option 4 over 5, or 5 over 4? To be sure, one would be loathe to perform either one, but are there facts that make choosing one option morally preferable to choosing the other? It would seem there are.

Option 4 reduces #2000's chances of survival to 0, whereas Option 5 reduces any one else's chances of survival by a fraction of one percent. On Option 5, no one's death is assured; each individual detainee's chances of survival remain effectively 1:2. On Option 4, #2000's death is assured. The argument then might be formulated this way. If morality (or a morally informed decision procedure) encourages us to choose that action the opposition to which is weakest, then we ought morally to choose Option 5 over Option 4, since the opposition to Option 5 rests on the fact that, relative to Option 4, it reduces any one detainee's chances of survival from 0.5002 to 0.4998—a difference of four-hundreds of one percent—whereas the opposition to Option 4 rests on the fact that it will result in #2000 losing his life.

These observations, however, support the following thought: even if we were committed to the claim that reasons for preventing harms are no more or less stringent than reasons against producing harms (and I am not), determinacy appears to play a role in considering our reasons against producing harms. On the one hand, we may not discriminate between Sadist X and Y as far as their motives are concerned, just as

we may not discriminate between those whose motives are to save the lives of others. On the other hand, were we in the regrettable position of having to deliberately harm others, discriminating between those acts that have as their objects determinate victims and those acts that have as their objects groups may well be morally appropriate.

III. Token Costs v. Type Costs

The third criterion in our list states that *S*'s executing those actions or contributing those resources can be expected to result in only a trivial cost to *S* when compared to the loss the victim faces. For this criterion to apply to the case of distant suffering, we must assume that my contributions to Oxfam can be expected to result in only a trivial decrease in my well-being when compared to the loss the victim faces. This assumption is difficult to make sense of in the case of distant suffering for reasons we just considered: there is no determinate individual at whom our assistance is directed. To speak of "the victim" here is deceiving. You are not sending a contribution to Oxfam to rescue any one individual; rather, you are hiring that agency to carry out activities it determines will improve the living conditions of those abroad. Unlike rendering aid to the drowning child or injured bicyclist, sending a monetary contribution to a relief agency is part of supporting a *collective good*, viz., humanitarian relief. But supporting a collective good bears little resemblance to the sorts of assistance one is asked to render in the case of the drowning child. But this fact has still other consequences we have yet to considered, particularly in relation to condition (c).

The fact of modern humanitarian assistance, in conjunction with the nature of collective goods, *can* yield substantial decreases in my well-being, even by comparison to the potential victim. Let me approach this point by way of example. Some months back I sent a modest contribution to the relief agency CARE. Less than two weeks later, I received a receipt acknowledging my contribution and thanking me for my generosity. But along with the receipt came another request for a donation. (Indeed, it was suggested that I arrange for an automatic monthly withdrawal from my bank.) Apparently, my contribution had not staunched the humanitarian crisis. Of course it would have been naïve to think it would. But I was not prepared for the slew of requests for assistance that suddenly appeared in my mailbox. Next to the third and fourth requests from CARE were requests from UNICEF, a local outreach program SOME (So Others Might Eat), a D.C.-area scholarship fund, Amnesty International, and the ACLU. (I now receive at least one solicitation from UNICEF or CARE *per week*.) It was hard to believe this was merely coincidence, and I found myself privately regretting my initial contribution. Indeed, as CARE itself recognizes, this was only the beginning.⁶² And as far as my role in humanitarian assistance is concerned, was there any end in sight?

David Schmitz (2002) distinguishes between *token-cost* and *type-cost* as a way of identifying the sorts of concerns I recently experienced. The former is just the cost to me of a particular rescue, whereas the type-cost is the cost to me of undertaking a kind of rescue whenever the occasion arises in my life. We can agree that the token-cost to me of sending a modest sum of money to CARE, like pulling

⁶² CARE's motto, "Where the End of Poverty Begins," might be read as a compendious way of saying "Where Ending Poverty Begins, But Never (Surely Not in Your Lifetime) Ends."

the child from the shallow pond, is trivial. But while the type-cost to me of rescuing individuals like the drowning child and the injured bicyclist will remain small (at least by any reasonable estimate), the type-cost of contributing to humanitarian relief efforts is not.⁶³

Tending to a local emergency today, one hardly expects to face one tomorrow. But this is precisely the situation one faces when attention is turned to humanitarian crises. After all, the fact that CARE encourages its donors to use monthly electronic-fund transfers is telling: These emergencies are not going away any time soon. They are, rather, “permanent feature[s] of your moral landscape.”⁶⁴ The upshot is that, while the token-costs would appear to be trivial in the case of distant suffering, the unlimited number of tokens yields a very high type-cost. It would seem, therefore, that condition (c) masks important differences between the case of distant suffering and the other cases we’ve been considering.

IV. Rescue and Commonsense Prohibitions

Finally, let’s consider condition (d), which states that my efforts at rescue under the circumstances do not require me to perform actions that are prohibited by commonsense morality. While initially plausible, this condition, too, belies problems. For example, it is possible to imagine “boring” cases that appear to conflict with commonsense morality. We must not forget that when I contribute to a collective good like humanitarian relief, I am depriving others of those resources, others to whom I may have an obligation to assist. Family members are, of course,

⁶³ See also Hooker (1999). McGuire (1985) offers a mathematically rigorous demonstration of how various changes in donor-recipient relations can yield sizable costs to agents.

⁶⁴ Schmitz (2000), p. 11.

the first to come to mind. Suppose my child is born with a minor, though conspicuous, disfigurement. When I decide to send the money I would otherwise spend on corrective surgery to Oxfam, does this count as a violation of commonsense morality? Arguably so.⁶⁵ But if it is, what distinguishes this case from other cases of bestowing largesse on family members? I would argue that commonsense morality contains a prohibition against placing the needs of strangers above the needs of one's family.

But aside from these more obvious cases, commonsense morality does not enjoin us to care for everyone equally. Indeed, tending to the needs of strangers over the needs of the members of your community can count as a moral failing—so, too, can tending to the needs of individuals abroad over the needs of fellow citizens. For example, a recent Census report indicates that 34.6 million Americans are currently living at or below the poverty line. Surely this fact is relevant in deciding on what I should do. Why? Because commonsense morality recognizes a special moral obligation to other citizens. All suffering should not be treated equally.

But my point can be made without having to establish these claims. For the very fact that such questions *can* be raised indicates a further place where the case of distant suffering diverges from our original cases. For were I to face a child's imminent drowning or the assault on someone like Kitty Genovese, it would be

⁶⁵ Consider the real-life case of Zell Kravinsky, who, after donating most of the family fortune to charity organizations, recently donated one his kidneys to a needy stranger ("An Organ Donor's Generosity Raises the Question of How Much is Too Much?" The New York Times (August 17, 2003): A12). But his altruistic campaign has alienated his parents, his children, his wife (who is threatening to divorce him), and his friends. According to the revisionist, Kravinsky's actions are worthy of the highest moral praise. But most people with whom I've discussed the case find his actions, apart from being mildly pathological, morally troubling. It is not at this point clear what the source of this uneasiness is, but this does not prevent us from saying that commonsense morality does not recognize a requirement to make such sacrifices.

patently irrelevant (not to mention, irresponsible) to wonder whether my efforts would be better spent elsewhere. But of course such questions are not at all irrelevant when I face a solicitation from Oxfam.

Development economists are deeply divided over just what I, as an individual, can do to improve the living conditions of the world's desperately needy. For example, is it so obvious that I ought to send my money to agencies that provide emergency relief rather than agencies that work to *prevent* the need for emergency relief? And what of organizations like Amnesty International that work for political reform? If, as some have suggested,⁶⁶ democracy is one of the most important factors in improving the living conditions of a nation's citizens, shouldn't my money go there? Perhaps, in light of deplorable conditions of displaced persons, I should be contributing to refugee assistance programs. Or, maybe I should be investing in corporations that invest in equity capital in developing countries, which would in turn promote a stable workforce and improved living conditions.

In the wake of ongoing humanitarian crises are very difficult questions about how I could best respond—so difficult, in fact, that the very people charged with generating answers cannot reach a consensus about what should be done. Moreover, nothing is said about other sorts of evils that we'd be better off without—e.g. racial injustice, ethnic cleansing, genital mutilation, slavery, terrorism, drug addiction, environmental destruction, and so and so on.⁶⁷ It would be hard to overemphasize how important this difference is in contrasting the case of distant suffering to the case

⁶⁶ Sen (1997).

⁶⁷ See Feldman (1999).

of the drowning child. For in the latter case, there is no question of how one ought to help.⁶⁸

If distant suffering is a case of easy rescue, then we're going to have to change the conditions constitutive of *ER*, for each of them failed to sufficiently capture important aspects of distant suffering. We saw that the case of distant suffering, in contrast to cases of local emergencies, involved a change in both the number and specificity of potential recipients, and this in turn forced us to confront the essential anonymity of those we wish to help, not to mention their "horizonless" numbers. Moreover, the interposition of a large bureaucracy introduced a form of unreliability that was missing in our original cases. It was for this reason that relieving distant suffering was more accurately characterized as a collective good, but this entails a fundamental change in the donor-recipient relationship and, it would appear, the content of our reason to provide assistance.

So why not alter *ER*? The problem with this of course is that altering *ER* to accommodate the case of distant suffering runs the risk of displacing our original cases, which, by most accounts, constitute paradigm cases of easy rescue. Without these the revisionist loses much of the force of his argument. So can the revisionist modify our original set of conditions (i.e. comprising type *ER*) in such a way as to make room for both the paradigm cases of easy rescue—viz., the drowning child, the

⁶⁸ Incidentally, these last remarks favor construing our contributions to collective goods as *imperfect* duties, since there would appear to be considerable latitude in both whom I should aid and how. By contrast, it would appear that I have a *perfect* duty to aid the drowning child, the injured bicyclist, Ms. Genovese, and so on, for under those specific circumstances, there is no latitude in what actions I ought to undertake. See Chapters 6 and 7 below.

injured bicyclist, Kitty Genovese, etc.—*and* the case of distant suffering? I believe that the answer is no. The revisionist, however, will accuse me of missing the point.

The revisionist sees no reason to alter *ER*, because while there is no doubt that the above differences separate the case of distant suffering from our original dataset, the point is that these differences are morally irrelevant. As far as deliberation is concerned, among the facts that should weigh heavily on our thinking are that facts referred to in *ER*. According to revisionists like Singer and Unger, that I can't be sure how or on whom my money will be spent should not make a moral difference. The fact that innocent people are in danger of suffering serious losses, combined with the fact that I can easily do something about those losses, are all the facts I need in order to know that I have moral obligation to render aid.

As tempting as it might be for the revisionist to lash himself to this point, this line of response is inadequate. First, it simply begs the question against the non-revisionist by dismissing all other concerns as morally irrelevant. I've argued, however, that considerations of consequence are not the only relevant considerations in determining the appropriate moral response. Moreover, in the absence of a fuller understanding of how and to what extent suffering provides those of us in a position to help with a reason to help, we can make no assumptions about how, if at all, we ought to assist those in need.

Second, it is apparent that the revisionist retreats to the truism about performing easy rescues and, taking cover behind the tremendous importance of easy rescue, refuses to be bothered by what he takes to be petty disparities. But the non-revisionist is not challenging the importance of easy rescue. The question pursued in

this chapter is whether distant suffering really is a case of easy rescue, as exemplified by *ER*. There is no inconsistency whatsoever in challenging the assumption that distant suffering is a case of easy rescue while at the same time emphasizing the moral significance of easy rescue as well as the despicable conditions in which millions of the world's inhabitants languish. Unger would have us think that challenging the former is merely a disguised way of challenging the latter.⁶⁹ But this is not so.

Finally, and in fairness to the revisionist, I have not yet offered a rival interpretation of our reasons to render aid. While I have argued that the consequentialist interpretation seems unable to account for important moral considerations, exactly what those considerations are and the extent to which they make serious moral demands of us I have yet to explain. I suggested that i) the relationship one stands in to a determinate aid-recipient represents an important moral dimension. Much more needs to be said before this idea can be considered part of a rival interpretation. Indeed, it is one thing to argue that there are differences separating the case of distant suffering from our paradigm cases; it is quite another to show that these differences can be accounted for by one's moral theory. For if the differences these aspects pick out cannot be traced back to the content of an acceptable moral theory, then we have what Kagan calls "dangling distinctions," that is, distinctions that "hang free" from the rest of one's moral theory in a way that leaves them unexplained and, hence, unjustified.⁷⁰ The challenge we must face in

⁶⁹ See Unger (1995), Chapter 11.

⁷⁰ Kagan (1989), p. 14.

later chapters involves showing that the differences we have identified are not left unaccounted for.

CHAPTER FOUR

Limiting Obligations

THE PREVIOUS TWO CHAPTERS consisted of various attacks on revisionism, the view that we are morally obligated to reduce or prevent suffering whenever and wherever it occurs. I argued that revisionism fails to capture much of the important content of our judgments concerning cases like the drowning child, and it is partly for this reason that we ought not assimilate this case to the case of humanitarian relief. In this chapter, I want to address another source of doubt about revisionism, viz., that it fails to recognize intuitive *limits* on what we can be expected to do for others. I have already expressed doubts about the over-demandingness charge (Chapter 2); however, there needs to be a thorough accounting of our intuition that we cannot be required to live lives of constant lifesaving.

There are two schools of thought on how such limits ought to be understood. One school of thought maintains that we ought to have the prerogative to act, at least sometimes, from what is called the “personal point of view.” From this point of view, my projects and relationships are said to take on a degree of significance out of proportion to their significance from the impersonal point of view. But since revisionism forbids such partiality, revisionism ought to be rejected. Another school of thought maintains that limits can be generated from within the “impersonal point of view” itself. On one such approach, it is noted that as a matter of contingent fact the level of compliance among potential donors is likely to be *greater* if donors are allowed to occasionally devote time and resources to personal projects that are of

special significance to the donor. Another approach in this school of thought maintains that revisionism without restriction (that is, one that conforms to an optimizing principle of beneficence) is unfair insofar as complying agents are required to “pick up the slack” for non-complying agents. These two approaches do not so much amount to a rejection of revisionism as a way of reining in its requirements.

While the arguments that take this line of attack are instructive in the way that they open up the range of reasons for beneficence, they tend to raise more questions than they answer. In particular, the emphasis on points of view obscures many of the important questions about beneficence, and almost no effort is made to connect “how things seem from *my* point of view” with reasons why I ought to identify with some projects over others—why, that is, I ought to identify with non-beneficent projects over beneficent projects. In sum, both schools of thought either fail to deliver the sorts of limits we seek or deliver the right limits but for the wrong reasons. At the same time, there is promise in the suggestion that noncompliance can set boundaries on our positive duties; I explore this possibility in Chapter 7.

I. The Argument from the Personal Point of View

Perhaps the most pervasive line of argument attempting to limit our obligations to reduce suffering is what I’ll call the argument from the personal point of view. In rough outline, the argument runs as follows. For any agent *A*, *A* stands in a special relationship to her projects, such that from her standpoint they possess a degree of significance out of proportion to the significance they possess from a completely

impartial standpoint. If *A* is always required to do everything she can to reduce suffering, then she is barred from treating her personal projects as having this greater degree of significance; she is unable, that is, to sometimes act from the personal point of view. But *A* ought to be able to sometimes act from the personal point of view. Hence, *A* is not required to do everything she can to reduce suffering. More generally, an acceptable system of moral obligations must not systematically prohibit the proper expression of *A*'s personal point of view.⁷¹ It follows that there will be conditions under which *A*'s decision to pursue her own projects, at the cost of not reducing further suffering, will not count as a moral wrong.

Despite my sympathy with the conclusion, I am skeptical that this line of argument can successfully reach it. The problem lies in sorting out the ambiguities surrounding the claim that agents ought to be able to act from the personal point of view. Most of the attention has centered around the question, on what grounds should an agent be able to act from the personal point of view? But there are other, more fundamental questions that need addressing: what is the function or role of “points of view” language in moral discourse? What does it mean to say that one has a reason to act *from* a point of view? The success of the argument from the personal point of view depends critically on how these questions are answered, for unless we know what the point of “points of view” language is, we cannot adequately assess the claim that we have a reason to adopt one point of view over another. Our inquiry, then, is

⁷¹ For example, Scheffler (1982: 57-8) argues that “we must surely reject any regulative principle for persons which ignores the independence of the personal point of view.” Thomas Nagel (1991: 15) agrees: “the personal standpoint must be taken into account directly in the justification of any ethical or political system in which humans can be expected to live by.” Even apologists for utilitarianism concede the significance of the personal point of view, but think utilitarianism possesses the resources to account for it: See, e.g., Kagan (1984); and Brink (1986).

shaped by answers to two questions: first, what is the function of points of view language in moral discourse, and second, given this function, what reasons do I—qua rational and reasonable agent—have for sometimes acting from the personal point of view?

I want to argue that how we answer the first question places important constraints on the content of our answer to the second question. To motivate this argument, I need to distinguish three different answers to our second question. According to what I'll call the *attachment* interpretation, we stand in a morally protected relationship to our freely chosen projects on the grounds that we care differently about our freely chosen projects, and such care can enter into our moral deliberations as such. Correlatively, the *personal integrity* interpretation contends that requiring that we sacrifice these freely chosen projects is a violation of our personal integrity, which is itself a moral violation. Finally, the *Kantian* interpretation maintains that moral agency must be sensitive not only to the happiness of others, but to one's own happiness as a feature of rational agency.

What unites these various approaches is the idea that abolishing our individual perspectives altogether would not move us closer to a moral utopia. But this is not because allowing agents to act from the personal point of view is, as Barbara Herman (2002) puts it, “a kind of minimum wage to keep moral workers happy so they won't go on strike.” Instead, our individual perspective, with its tendency to give greater weight to personal projects, constitutes an independent source of *moral* value. It is not, however, always easy to see why. Part of the problem, as I'll suggest shortly, is that these answers fail to take full account of the *function* of points of view language.

The Attachment Interpretation. The attachment interpretation attempts to establish boundaries on what we are required to do for others by way of pointing to the special attachment we feel towards our freely chosen projects. As natural as this suggestion might seem, I want to draw attention to the difficulties it faces.

The first thing to note is that there is little, if any, space between the claim that ‘*M* is the set of my projects’ and the claim that ‘I am closely identified with *M*.’ That I care differently about *M*, that I am especially concerned with *M*’s fate, that *M* has a distinctive claim on my attention, are all more or less trivial consequences of the fact that *M* happens to be the set of my freely chosen projects. Nor is it any wonder that, as Bernard Williams (1973: 116) puts it, I should be identified with my actions “as flowing from” *M*. But if this is so, then of course it is no surprise that I stand in a special relationship to my projects, since this is effectively what having projects consists in. But now the question arises: why should this fact be the source of special moral consideration? Since my relationship to my projects and the distinctive claim they have on my attention follow more or less trivially from the fact that they are *my* projects, and since it is implausible to think that a trivial truth could be the source of moral concern, how does the special relationship I bear to my projects derive its special moral significance? Even if it is true, as Scheffler (1982: 57) claims, that an agent “cares differently about his projects just because they are *his* projects,” it’s far from clear why this should be assigned special moral status.⁷²

⁷² Scheffler continues:

His interest in how they fare is not ordinarily exhausted by estimating the value or disvalue from the impersonal standpoint of their success or failure, and feeding this estimate into the impersonal calculus in order to arrive at an overall assessment *sub specie aeternitatis* of a

Surely there must be more to the claim than this. To see that there is, consider how “points of view” language has been used in the past—not to describe the relation between an agent and his projects, but, roughly, as a method for determining *which* projects he ought to adopt. Kant, for instance, suggests that it is only when we “abstract away from the personal differences of rational beings and also from the content of their private ends” that the “kingdom of ends” comes into view.⁷³ Mill, in recommending the strict impartiality of “a disinterested and benevolent spectator,” attempts to balance the significance of one’s projects against the projects of everyone else.⁷⁴

More recently, Thomas Nagel (1986) has attempted to give a systematic account of how these various points of view function. For Nagel, taking up the impersonal point of view involves forming “a conception of the world as centerless—as containing ourselves and other beings with particular points of view;” this is supposed to open up the possibility of recognizing “values and reasons that are independent of one’s personal perspective and have force for anyone who can view the world impersonally.”⁷⁵ But, as Nagel himself accedes, a completely “centerless” conception of the world is impossible for creatures like us.⁷⁶ Try as we may, a complete transcendence of our place in the world is out of reach. But the point is in the process: we are seeking, as far as possible, to “occupy a standpoint detached from

contemplated outcome. His own projects and commitments have a distinctive claim on his attention: he cares about them out of proportion to the relative weight carried in the impersonal calculus by his having and caring about them.

⁷³ Kant (1996): 4:443.

⁷⁴ Mill (1863/2001): p. 17. Hume, too, was not averse to such talk. For example, he distinguishes *moral* goodness or evil from other sorts of evaluative notions by the way we consider another’s character “in general, without reference to our particular interest:” see Hume (1978): Bk. III, pt. 1, sec. 2: p. 472; see also Bk. III, pt. III, sec. 1: p. 582.

⁷⁵ Nagel (1986), p. 140.

⁷⁶ See also Hare (1981): Chapter 3; and Friedman (1989).

that of our purely personal desires and interests,” for it is from this standpoint that we can “correct inclination and discern what we really should do.”⁷⁷ The central function, then, of the impersonal point of view is to mitigate the distorting effects of how things seem from our individual points of view. And this coincides with the historical examples above.

So, generalizing, we might say that adopting a “point of view” is supposed to be a way of determining whether or not an agent is allowed to make reference to herself in assigning significance to various ends or projects. From the personal point of view, she is; from the impersonal point of view, she is not. If this is right, then the distinction of interest does not turn on differing points of view per se, but on the *scope* of what one takes one’s projects to be, or the reasons one has for pursuing the projects one pursues. This would cohere with Nagel’s suggestion that the aim is to “correct” the distorting effects of self-interest.

But now it appears that the mere fact that my relationship to my projects is unique—they’re my projects, after all—will be true *no matter what my projects are*. Emphasis on attachment, in other words, misses the more fundamental question: which projects should I pursue, and why? It is tempting to think that in the event that I cease to assign special significance to a subset of all projects (those I call “mine”), instead treating all projects as equally significant, I have achieved the impersonal point of view. This would be a mistake. For what I would have achieved, were I capable of such a thing, would be an expansion of those projects I call “mine.” It is not that I no longer give partial treatment to “my projects;” it is that I have widened—to the fullest extent possible—the scope of my projects. I see a reason to treat all

⁷⁷ Nagel (1986), pp. 8 and 140, respectively.

projects as having value. But then it becomes an open question whether such an attitude is, if possible, desirable.⁷⁸ Again, it's unclear how the appeal to the personal point of view—when this is simply a way of describing the special attachment we have to our projects—is supposed to supply the needed moral content in virtue of which our obligations to reduce suffering can be limited.

The same point can be seen made from another angle. Suppose that when I adopt the “impersonal point of view,” I identify with all projects; they all have a distinctive claim on my attention—distinct, that is, from any particular project. And when I adopt the “personal point of view,” I identify with only a subset of all projects. What we see is that in both cases there is reference to a given agent (me) and reasons (I have or should have) for pursuing those projects. But, again, when we take this view of the distinction, the importance of the personal point of view is diminished. For whatever is interesting about the fact that I bear a special relationship to my projects will be true no matter what the scope of my projects are.

What is more, by locating the real source of the distinction in questions about which projects I ought to identify with, we effect a change in the sorts of considerations that bear on the question of which projects I ought to identify with. For example, among the considerations I might cite as reasons for valuing some project *t*, I might cite the fact that *t* is exciting or challenging or beneficent or pleasing or some such thing. What I would not cite, however, is the fact that *t* is *my* project. Such a reason for valuing *t* would be, at best, self-aggrandizing. At worst, it would be circular: I identify with *t* because I identify with *t*. Reasons for valuing *t* make

⁷⁸ On this question much has been written. See especially Hampton (1993); Wolf (1982); Cottingham (1986); and Walker (1991).

reference to *t*'s various qualities. That I am attached to my projects is not an additional reason for valuing *t*; it is rather a consequence of taking other qualities of *t* as reasons for pursuing and sustaining *t*.

The Personal Integrity Interpretation. But perhaps I am moving too fast. There is, according to the personal integrity interpretation, a sense in which my attachment to *t* is a reason for valuing *t*: when I am asked to sacrifice *t* for, say, the general good. More specifically, that I identify with *t* is a reason for treating my sacrificing *t*, after I have *already* formed an attachment to *t*, as a violation of an important moral virtue, viz., personal integrity. On this interpretation, the claim that our individual perspective constitutes an independent source of moral value should be understood as a claim about what it means for an agent to have to sever his attachment to his projects. Williams (1973: 116-7) puts this point this way: Demanding that an agent always deliberate from the impersonal point of view “alienate[s] him in a real sense from his action and the source of his action in his own convictions.” But, urges Williams (ibid), “*his* actions and *his* decisions have to be seen as the actions and decisions which flow from the projects and attitudes with which he is most closely identified. It is thus, in the most literal sense, an attack on his integrity.”

But suppose we were correct above in thinking that the distinction of interest is *not* in points of view per se, but in the scope of the projects with which agents identify. Then an agent who (perhaps per impossibile) closely identifies with all projects will be forced to violate her personal integrity if she is asked to identify with only a narrow subset of those projects. So the question becomes: does the personal

integrity interpretation provide any sort of guidance as to what the scope of my projects ought to be? It would appear that the only thing the personal integrity interpretation recommends is that once the scope of my projects is fixed, once I have come to identify with a given range of projects, then I ought to exercise my moral autonomy and resist the requirement that I constrict or expand that range. In this way, my personal integrity is preserved. But on the question of taking on the project of reducing suffering, the personal integrity interpretation, like the attachment interpretation, is silent. If, that is, the claim that agents ought to be able to sometimes act from the personal point of view is to be understood as a claim about the sorts of reasons we have for adopting some projects over others, then the personal integrity interpretation is of no assistance.⁷⁹

It might be responded, however, that, for most of us, taking on the project of reducing suffering would require substantial sacrifices in those projects we *already* identify with. After all, the question of which projects we ought to identify with is asked *in media res*. So while the personal integrity interpretation may not specify which projects one ought to adopt, the fact that most of us have already adopted projects that would require sacrificing if we undertook the project of reducing suffering provides sufficient content to the claim that agents ought to sometimes be able to act from the personal point of view.

⁷⁹ The personal integrity interpretation faces an additional problem: if personal integrity amounts to a sort of moral immunity, on what grounds can individuals be restricted from pursuing projects that needlessly harm other individuals? If the personal integrity interpretation is indeed silent on the question of which projects one ought to pursue, this worry should be especially troubling. For Scheffler (1983), there are no foundations for deontic restrictions, so this result is no worse than the alternative, whereby restrictions are given no theoretical backing. Others, myself included, express serious reservations about granting individuals the sort of immunity entailed by these interpretations: e.g., Kamm (1996), Chapter 7; and Myers (1999), Chapter 2.

There can be no doubt that for most of us taking on the project of reducing suffering will require of us substantial sacrifices, but, as I argued in Chapter 2, the question of the “demandingness” of such a project cannot simply be assumed. This response assumes that we are *entitled* to those projects to which we have committed ourselves, but such an assumption cannot go undefended. Perhaps such a defense is in the offing. But the bare notion of personal integrity seems an unlikely source for such a defense. I would hardly think, that is, that defenders of the personal integrity interpretation would want to explain entitlement solely in terms of what would be costly for an agent to sacrifice. For then art thieves would be entitled to their stolen masterpieces and nephews would be entitled to murder their rich uncles.

The dilemma facing the attachment interpretation and the personal integrity interpretation can be summarized as follows. If our reason to sometimes act from the personal point of view is essentially a reason to sometimes devote greater energy and resources to those projects with which one identifies, then it is virtually a trivial truth that we have such a reason since this is part of what it means to see certain projects as one’s own. But if our reason to sometimes act from the personal point of view is essentially a reason to adopt a set—presumably, a narrow subset—of projects with which to identify, then the two interpretations under consideration leave the content of that reason unspecified.

The Kantian Interpretation. As a way out of the dilemma, we might turn to the Kantian interpretation. On this interpretation, the reason we have to sometimes act from the personal point of view is essentially a reason to adopt a set of projects that,

as Herman (2002: 238) puts it, draws “own-happiness into the space of moral reasons.” Drawing on Kant’s account of obligatory ends and the imperfect duties they give rise to, Herman (ibid: 242-3) defends the idea that perfecting one’s own person is, inter alia, a precondition for effective moral agency:

[U]nless one is willing and to some degree able to enjoy life, one cannot appreciate and so correctly evaluate the range of human concerns. One will not make wise judgments about either one’s own needs as an agent, or about the happiness of others...The conditions for effective moral agency are not, then, to be regarded as luxuries. Although to an extent contextually specific, and very often resource-demanding, the cultural conditions of moral agency are matters of moral necessity.

In the language of “points of view,” we might summarize Herman as saying that unless we are allowed to sometimes act from the personal point of view, we will fail to appreciate not only the appropriate means of assisting others, but the means of recognizing those who stand in need of assistance. For it is only when we attend to our own needs and the appropriate means of satisfying them that we can, from the impersonal level, recognize need and what to do about it. Unlike the two interpretations we considered above, the Kantian interpretation does make recommendations about the sorts of projects with which we should identify: at least some of these projects should promote our own happiness without regard to the happiness of others, since, according to Herman (ibid), “the positive experiences of free enjoyment enable moral judgment.”

But in response to this suggestion, we might reasonably wonder: How far must our moral development extend before we can effectively engage in beneficent activities? Is moral development a lifelong pursuit, such that there will always exist a

“moral necessity” for developing the conditions of our individual moral agencies? If so, then it might seem that I can put off the project of assisting others indefinitely, since the preconditions of effective moral agency are never quite met. But if, on the other hand, I can meet the preconditions for effective agency early on as it seems I can, then it might seem that I can dispense with the personal point of view altogether, since its purpose has been fulfilled. After all, even young adults can understand, if not fully appreciate, the seriousness of living under conditions of extreme poverty and what they themselves can do to reduce such poverty.

This response, however, does not take full account of the role of moral development. A moral agent seeks to perfect her capacity to respond not only to the needs of others, but to her own needs as well, since, qua rational agent, she must treat the humanity in herself as an end and not merely as a means. In exercising one’s beneficence, therefore, one must not ignore the duties to oneself. Among the obligatory ends Kant identifies (and Herman endorses) is the end of perfecting one’s own humanity, where this is traditionally understood as an agent’s capacity to set ends. At the same time, however, we have as an end the happiness of others. But Kant (1996) famously allowed that obligatory ends leave “a playroom (*latitudo*) for free choice in following...the [moral] law.” But Herman (2002: 238) does not view this as simply a license to sometimes forgo assisting others or, for that matter, ignoring one’s own needs; rather, “it is only a permission to limit one maxim of duty by another.” Hence, according to Herman (*ibid.*), the “obligatory end of one’s own perfection [can sometimes limit] action for the end of others’ happiness.”

On this interpretation, then, it would be a mistake to see our relation to our projects as beyond the reach of moral requirements. Morality can dictate that some of the projects we promote serve no other function than furthering the natural interest we have in our own happiness. If this is right, then sometimes acting from the personal point of view is not merely a morally permissible alternative, but a moral requirement.

Of the interpretations we've considered thus far, the Kantian appears to be the most promising, since it goes beyond the uninteresting idea that one should take a particular attitude towards one's projects or the imperative to alter them. It stipulates a condition on the adoption or maintenance of one's projects: at least some of them should promote one's own happiness. Of course there will be individuals who, by nature, find satisfaction in nothing less than devoting themselves to charitable works. But this possibility does not threaten the interpretation under consideration, since it is enough that some of those projects are enjoyed for their own sake.

The claim that morality requires that some of the projects we adopt promote own-happiness admits of (at least) two explanations. Unfortunately, neither is satisfying. On the first explanation, the reason some of the projects we adopt should promote own-happiness is because denying own-happiness reduces one's willingness or capacity to aid others. On the second explanation, the reason is, morality requires the development of one's capacity to experience own-happiness since this is entailed by a more general obligation to perfect oneself.

Herman resists the first explanation, as well she should. For the value of own-happiness is not simply a function of its usefulness in promoting the happiness of

others. If it were, then we would each have a reason to develop, as far as possible, lifestyles of greater and greater asceticism. Since there is no intrinsic value in own-happiness, then the less we must rely on it (to perform acts of beneficence) the better. But this is the wrong way to grasp own-happiness. Surely part of what makes life worth living is the ability to enjoy it, and this constitutes a reason to enjoy it—whatever else morality has to say about the matter. There is, however, a more serious problem with this first explanation: to deny that own-happiness is intrinsically important while claiming that other's happiness is leads to the bizarre idea that others have a reason to protect or promote my happiness but I don't.⁸⁰ But why should my happiness count only when it is viewed "from the outside?" With results like these, the first explanation ought to be rejected.

The second explanation, however, goes too far in the other direction. My lack of interest in my own happiness might be considered pathological or unusually self-abdicating, but it is a further claim to say that I act *wrongly* in failing to take an interest in my own happiness. This is not to deny that the ability to enjoy life is partly constitutive of a life worth living. The question concerns the reason(s) we have for cultivating this capacity: is it a distinctly moral reason? I do not see that it is. Suppose I do not take an active interest in my own happiness, instead turning my attention to the needs of others. And suppose my interest in others is not simply a means of gratifying myself; I simply view the serious suffering of others as a compelling reason to devote my life to reducing their suffering. Or perhaps it is a reflection of religious devotion. In any event, I consider my needs met. But if it is true that I have a moral obligation to take an interest in my own happiness, then,

⁸⁰ See Slote (1985), Chapter 1; and Darwall (1983), Chapter 15.

ceteris paribus, we would have to say that my life is of less moral value than my twin who attempts to enjoy those same charitable works. But unless we're prepared to revert back to the first explanation, there seems no reason to accept this. Why should the difference necessitate a moral diagnosis? Why is it not enough to say that the difference is one of psychology? I might be criticized for missing out on an important aspect of life; it might even be said that my life is not worth living. But surely that is an evaluation only I can make. And if I respond that my projects give my life meaning, it strikes me as paternalistic, at best, to say that my attitude nevertheless represents a moral failing.

Now such lack of interest should not to be confused with cases of *active* self-abdication, where I actively suppress an interest in my own happiness for, say, the sake of others.⁸¹ There, my actions may in fact be open to moral criticism because I treat myself merely as a means. But while it may be true that morality forbids me from treating my own capacities and talents as dispensable goods, it is a further claim to say that morality requires that I develop my capacity to "enjoy life."

It should also be noted that the role of happiness in Kant's ethics is a highly disputed matter, and Herman's reading of Kant glosses over important difficulties. For example, Kant's official position with respect to happiness more closely resembles the first of the two reasons discussed above, viz., that without some attention to one's own happiness, one will be less likely to conform one's behavior to morality's dictates. But, as Herman herself acknowledges, this is a crude way of understanding the role own-happiness plays in our lives.

⁸¹ See Hampton (1993)

It is, instead, Kant's discussion of meritorious duties to oneself that Herman cites in support of her account. The relevant passage in the Groundwork reads: "Now there are in humanity capacities for greater perfection which form part of nature's purpose for humanity in our person. To neglect these can admittedly be compatible with the maintenance of humanity as an end in itself, but not with the promotion of this end."⁸² Kant's claim is that among the goals or purposes inherent in nature is the perfection of our natural talents and capacities, among them (presumably) the capacity to enjoy life. It is, in short, part of the function or purpose of humanity that we should perfect it.

But apart from Kant's own defense of this claim (which is elliptical, at best), it seems unlikely that a satisfying defense is forthcoming. After all, it would be a mistake to appeal to the biological sciences for evidence, since the teleological explanations on offer there do not license the sorts of normative claims Kant needs. Moreover, for reasons adduced above, it is not obvious that those who are disposed to take an interest in the needs of others over their own happiness are defective. While I may not choose such a life for myself, it seems unduly presumptuous to think that someone who does is "malfunctioning."

II. Moral and Nonmoral Reasons and the Pervasiveness of Morality

Both explanations of the claim that some of the projects we adopt should promote own happiness assume that own-happiness should be drawn into the space of moral reasons. But one suggestion we have yet to consider is that own-happiness—or the

⁸² Kant (1996), 4:430.

sorts of considerations that count in favor of it—functions as a limiting condition on morality.⁸³ On this view, reasons to protect, if not promote, own-happiness operate outside the moral domain, acting as injunctions against the infringement on personal projects and commitments. If this is right, then morality does not pervade all aspects of our lives. More specifically, this means that within our practical lives there is a domain in which moral assessment is not appropriate or does not apply. Some actions, attitudes, reasons, etc. are beyond the reach of moral assessment: they are neither morally justified nor morally unjustified.

It might be argued, then, that there is an important reason for allowing agents to sometimes act from the personal point of view, but this reason does not rest on the alleged moral importance of such a prerogative. Instead, the reason makes reference to considerations outside the moral domain—for example, that we have a reason to lead worthwhile lives. A life without the prerogative to sometimes act from the personal point of view, it might be argued, is not a life worth living. Leading worthwhile lives, then, functions as a consideration that can limit the demands morality can place on us. Moreover, this way of interpreting of the claim naturally sets constraints on the sorts of projects one should adopt: the projects one takes on should not make one's life not worth living.

Though initially attractive, this way of defending the personal point of view argument is less so on reflection. Apart from the difficulties associated with drawing

⁸³ There may be a way of reading the personal integrity interpretation in this way. Williams's (1972) first remarks on the subject suggest that a threat to personal integrity represents a moral consideration whose weight in moral deliberations is non-negligible. Later, however, Williams moves in the direction discussed in the text: reasons of personal integrity do not stand in need of moral justification, since they are beyond the domain of moral adjudication; see Williams (1985). For skepticism about this position, see Scheffler (1992): Chapter 1.

a principled distinction between moral and nonmoral considerations,⁸⁴ there is little reason to accept the thesis that some reasons, actions, attitudes, etc. are beyond moral assessment. Of course one might be drawn to this thesis if its negation implied that for every action, attitude, reason, etc., one must engage in a process whereby the moral status of the relevant token is determined. But the negation of the thesis does not imply such a result. It implies only that it is, in principle, possible to evaluate the moral status of any token action or attitude. Whether one should engage in such a process is a separate question, as is the question of how interesting the results might be.

Consider my act of making a pot of coffee. It might be argued that such a trivial action is beyond the reach of moral evaluation. This is not to say that the action is morally justified, but uninterestingly so. It is to say that for actions of this sort moral justification is simply inapplicable. But what motivates such a claim? It cannot be the fact that seeking a moral justification for the action is morally inappropriate or indicative of a troubling compulsion, since what is at issue is not *the act of seeking* a justification. Neither can it be the fact that under no circumstances is it conceivable that making a pot of coffee is morally impermissible. When, instead of calling the police to report a sexual assault taking place outside my window, I proceed to make a pot of coffee, it is quite natural to say that I should not have made a pot of coffee. I should have called the police! Alternatively, it would be wrong of me to make a pot of coffee if I promised not to—because, say, the electrical surge would threaten to blow a fuse and disrupt my colleagues computer work. Any number of cases could make the same point.

⁸⁴ See Raz (1996).

So it must be that, *under standard circumstances*, making a pot of coffee is not open to moral evaluation. But conceding this amounts to conceding that there are circumstances under which actions of this sort are open to moral evaluation. Why, then, should we say that when circumstances are standard, the possibility of moral justification is blocked? Surely it is more natural to say that under standard circumstances actions like making a pot of coffee are obviously morally justified. Just because an action is patently morally permissible (or, for that matter, impermissible) does not render the question of its permissibility inappropriate or out of place—uninteresting, perhaps, but not inappropriate.

If this is right, then a defender of the argument from the personal point of view cannot claim that occasionally acting from the personal point of view is beyond moral evaluation altogether, and so can place constraints on the demands of morality. At most, she can claim that so acting is morally justified, but uninterestingly so. But this is to beg the question with prejudice. The question we've been examining is precisely what the grounds, if any, are for claiming that agents ought to sometimes be able to act from the personal point of view. The suggestion under consideration effectively changes the subject: there need be no grounds for such a prerogative. But nothing so far recommends such a view.

But one might still argue that nonmoral considerations can “override” moral considerations since, as Raz (1996) argues, moral reasons do not possess a special kind of stringency. According to Raz, moral considerations do not make up a special class of reasons. As an example of how nonmoral considerations can override moral ones, Raz offers us the case of a young woman who decides to give up her job

assisting the homeless to have a baby. Raz stresses that at the time of her decision, there are no moral reasons for her to have a baby while there *are* moral reasons for her to continue working for the homeless. But if we judge that she does no wrong in giving up her job and having a baby, we must concede that there are cases in which moral reasons can be overridden, or “at least neutralized,” by other, nonmoral, considerations. If this is right, then it might be claimed that the considerations arising from the personal point of view are not only nonmoral, but are of the sort that can override moral reasons. This, in turn, would license the claim that one should be able to sometimes act from the personal point of view.

I believe there is much to recommend this line of thinking insofar as it takes seriously the variety of reasons that apply to us in our practical lives. My own way of approaching the issue of beneficence is by way of such reasons, where these do not form a cohesive set. But it should be born in mind that this is not itself an endorsement of the argument from personal point of view. After all, it is one thing to emphasize the heterogeneous nature of reasons to rescue, quite another to suggest that all (or any) of these reasons flow from the personal point of view. Indeed, it strikes me that the class of nonmoral reasons that override moral reasons will be small, so small as to be unhelpful in fixing on limits.

The argument from the personal point of view contends that it is in virtue of some special feature of our individual perspectives that our obligations to reduce suffering are limited. We have considered a variety of ways this claim might be interpreted, and all of them have been found wanting. It may turn out of course that a closer examination of practical reason—both its content and relation to action—may

prompt us to revisit the notion of the personal point of view. Until that time, I will set aside the argument from the personal point of view.

III. Limitations from the Impersonal Point of View

The aim of the argument from the personal point of view was to limit our obligations to reduce suffering, to occasionally free agents from the burden of having to reduce suffering. It might be thought that such limitations are only available from within the personal standpoint. In this section, we consider two approaches that attempt to yield limitations from within the impersonal standpoint. On the first approach, principles of beneficence must be tailored to the particular psychologies of complying agents, a constraint that can be generated without reference to any particular point of view, i.e. from the impersonal point of view. On the second approach, principles of beneficence must take into account the level of compliance within existing populations and the extent to which this places unfair demands on individuals that comply. Like the first approach, this approach does not rely on the personal point of view. While interesting lessons will emerge in the course of our critique, I will eventually reject both of these approaches as originally conceived.

“Workable” Principles of Beneficence. It is sometimes pointed out that a moral theory that demands that agents devote their lives to reducing suffering is unlikely to have the desired effect—indeed, it may produce just the opposite effect. The reason for this is not hard to see: as creatures prone to selfish desires we cannot, as a matter of contingent fact, be expected to live lives of extreme beneficence. And invoking a

morality that makes such demands has the effect of placing moral decency out of reach for most of us. The result may well be a decrease in compliance, since agents will wonder what the point of making any effort at all is when even very strenuous efforts to reduce suffering do not meet the minimum standard of decency.

One way of avoiding this outcome is tailoring principles of beneficence to individual capacities. R.M. Hare (1981) has developed a two-level account of utilitarianism that appears to account for such variation.⁸⁵ On this view, the prima facie principles we as individuals ought to adopt should be compatible with our individual capabilities and overall characters. In Hare's words, the principles should be "practically workable." And determining which ones are practically workable requires attending to one's character and circumstances:

if I were very saintly like Albert Schweitzer or Mother Theresa, I might have much more exacting principles than in fact I have; and the same would apply in the case of conduct in the face of danger if I were more courageous than I am. This last case illustrates very well the point that principles have to be suited to those who are to follow them...Each of us, therefore, has to ask himself what is the level of saintliness of which he is likely to be capable, and strive for that (Hare, 1981, pp. 200-1).

My moral obligations to prevent suffering, then, cannot outstrip my ability to discharge them, *given the kind of person that I am or am likely to be*. But this not simply the familiar lesson of ought-implies-can; rather, it is an outgrowth of Hare's two-level system of moral thinking.

Hare distinguishes between an *intuitive* level of moral thinking, which involves a reliance on prima facie principles, intuitions, and general dispositions in

⁸⁵ Peter Singer endorses this view in his (2004).

deciding how to act, and a *critical* level of moral thinking, which involves evaluating and ultimately choosing principles according to an independent set of theoretical constraints. Because we lack both the time and processing power to consider all possible outcomes various actions may yield, the intuitive level of thinking is, according to Hare, indispensable. But because the principles, intuitions, etc., at the intuitive level stand in need of justification, the critical level is likewise indispensable. Together, these two levels of moral thinking can generate the desired limitations in the following way.

From the critical level it is clear that principles enjoining agents always to promote others' welfare and never their own will not produce as much good as an alternative principle, one that allows agents to sometimes promote their own welfare. I cited several reasons for this above. Hare (*ibid*: 202) cites the fact that I cannot always be certain what would promote the interests of others, whereas I can be more or less certain what would promote my own interests. If more utility can be realized through principles enjoining agents to sometimes act for their own happiness, then we should work toward establishing just those *prima facie* principles—"backed up by powerful moral feelings, and attached to rather general characteristics of actions and situations"—at the intuitive level of thinking, since this is the level from which creatures like us must do our moral thinking.⁸⁶

⁸⁶ Hare (1981), p. 59. Hare (1964; Chapter 4) uses the analogy of driving a car. When we drive, we pay little attention to the principles we may have internalized over the course of our driving experience. We respond quickly and almost thoughtlessly to changing conditions, which suggests that most of us (thankfully) have received adequate driving training. Extending this idea to the moral realm, Hare suggests that much, if not all, of our present moral experience is a reflection of the sorts of moral training we received in the past.

Now those who, like myself, reject consequentialism as an overall moral framework will be tempted to dismiss Hare's suggestion without further ado. But I believe this account fails on its own terms. First, it's hard to take seriously the idea that we cannot be certain what would promote the interests of individuals threatened by, say, starvation. Ambiguity may surround issues like illiteracy or even chronic poverty, but it seems dangerously pedantic to plead ignorance in the face of mass starvations or epidemics.

But this worry aside, Hare's treatment of own-happiness falls into the same confusion that arose in the discussion of Kant. Although committing ourselves to various personal pursuits may serve the purpose of increasing the level of beneficent activity, it is a mistake to say that that is the reason for committing ourselves to such pursuits. Our reason for committing ourselves to such pursuits rests on, *inter alia*, the reason we have for living a worthwhile life or engaging in valuable pursuits. And the value of living a worthwhile life or engaging in valuable pursuits is not a function of maximizing utility. We have a reason for such things *even if* there are no moral reasons for such pursuits. This is not necessarily an attack on consequentialism. After all, consequentialism must commit itself to one nonmoral value or another if it is to have something to promote. Instead, the problem has to do with the content of the reason we have to sometimes promote our own happiness. Hare's account, I claim, assigns the wrong content to this reason.

Finally, the suggestion that I should ask myself what level of saintliness I am likely to be capable of and strive for that underdetermines the sorts of standards to which I ought to hold myself. Part of the problem lies in thinking that I can know, or

even be reasonably be sure of, the level of saintliness I am capable of *before* engaging in “saintly” activities. To say that “you won’t know until you try” is not ungrounded optimism; it’s a reasonable hypothesis.

Perhaps the most troubling aspect of Hare’s proposal is his view that moral principles be idiosyncratic. If individual capabilities and overall characteristics can alter the reach of principles of beneficence, why not think they can alter the reach of, say, principles of non-interference? If I am capable of containing my aggressive behavior only on occasion, then shouldn’t I strive to contain my aggression *only* on those occasions? If so, then I do not violate principles prohibiting aggressive behavior on those occasions when I cannot restrain my aggressive behavior since, on Hare’s proposal, there are no such principles. I am only bound to those principles that are tailored to my capabilities. This is an awkward result to say the least. Note that this is a different objection than the one consequentialists often face, viz., that it cannot generate side-constraints or (something approaching) deontic restrictions. Instead, this result threatens any view permitting the tailoring of principles to individual capabilities in the way that Hare suggests.

Compliance and Fairness. An alternative attempt⁸⁷ to generate—from the impersonal point of view—limits on our obligations to give begins by noting the (contingent) fact that few individuals do much to reduce global suffering. The relevance of this fact becomes clear when we consider the demands this places on an individual who seeks to act on her reason to reduce suffering. Because there is so little effort on the part of the general population to reduce suffering, the scope of that suffering is effectively

⁸⁷ Murphy (2000).

“horizonless.” Hence, the individual who intends to act on her reason to reduce suffering will have to make significant sacrifices in her current level of well-being. For as long as she is prepared to act on this reason, and as long as most other people are not, then horizonless suffering will require horizonless beneficence (assuming the absence of any agent-relative reasons limiting beneficence).

According to this line of thinking, what generates unacceptable demands on complying agents is the lack of compliance on the part of others. If most people did comply with general principles of charity, not only would the scope of suffering be narrowed, but the level of sacrifice for any one complying agent would be reduced. As it stands, complying agents have to “take up the slack” for those who do not do their part. And this, it is argued, is unfair. Thus principles of beneficence ought to be sensitive to facts of compliance since such facts are responsible for the unacceptable demands placed on complying agents. One way of capturing this suggestion is to allow that the requirements of beneficence at full compliance provide a benchmark for the determination of demands under partial compliance. Doing my fair share consists of doing what would be required of me *when everyone does her fair share*. Partial compliance cannot generate added burdens.

Murphy (2000: 80) formulizes this suggestion with the following “compliance condition:”

Agent-neutral principles should not under partial compliance require sacrifice of an agent where the total compliance effect on her, taking that sacrifice into

account, would be worse than it would be (all other aspects of her situation remaining the same) under full compliance from now on.

The rough idea is this: I should not be required to sacrifice more than I would have to sacrifice under full compliance. The reason, as he makes plain elsewhere, is that as the distribution of responsibility under full compliance should be fair so should the distribution of responsibility under partial compliance. Murphy sees this is an argument against what he calls the “optimizing principle of beneficence,” according to which agents are always required to do as much as possible for others—which is effectively what I have been calling revisionism.

Now the compliance condition is not itself a principle of beneficence; it is, rather, a constraint on any acceptable principle of beneficence. But Murphy does offer a principle of beneficence, one that combines the consequentialist principle of utility-maximization with the compliance condition. The “collective principle of beneficence,” as he calls it, runs as follows: “a person need never sacrifice so much that he would end up less well-off than he would be under full compliance from now on, but within that constraint he must do as much good as possible” (Murphy 2000: 86-7). The principle is meant to express both the demands of morality and its limits. What is significant, for our purposes, is the fact that these limits are supposed to come from within morality itself. For the compliance condition is grounded on considerations of fairness. Moreover, the importance of fairness is supposed to be independent of how things appear from the personal point of view.⁸⁸ That it would be

⁸⁸ But can considerations of fairness, on which the compliance condition rests, be defended on consequentialist grounds? Paul Hurley (2003) argues persuasively that the answer is ‘no.’

unfair to require that I do more under partial compliance (than under full compliance) can be determined without reference to my projects, preferences, desires, and so. In this regard, Murphy's proposal marks out an interesting alternative to the approaches considered in previous sections.

Apart from its role in the collective principle of beneficence, the compliance condition is also meant to provide a plausible diagnosis of the problems with revisionism, the view grounded on the optimizing principle of beneficence. Recall from Chapter 2 that despite the sense that revisionism was unacceptably demanding, no adequate explanation of this sense could be provided. To show that the problem of over-demandingness was a legitimate problem, one needed to show that there was something about costly *positive* demands in particular that were objectionable. That a moral principle in general was costly was not a sufficient reason for rejecting that principle. But, argues Murphy, the compliance condition can show that over-demandingness is a legitimate problem—and, moreover, that revisionism is over-demanding. On this reading, the over-demandingness charge is essentially a charge of unfairness: demanding that agents always do as much as they can to reduce suffering is unfair, since, under nonideal conditions (i.e. partial compliance), complying agents must take up the slack for non-complying agents. And this is essentially demanding that agents do more than their fair share. As Murphy sees it, the “real problem with the optimizing principle of beneficence is its unfair imposition of responsibility rather than the extent of its demands” (p. 99).

On the one hand, the compliance condition rightly captures the intuition that assisting strangers in need, particularly in the context of distant suffering, is and

should be a cooperative enterprise.⁸⁹ The burden of addressing distant suffering should be shared collectively, and in this context, no one should have to do more than his fair share. On the other hand, the compliance condition, as Murphy presents it, is open to obvious counterexample. For it is one thing to claim that, at full compliance, I should not be required to do more than what others are doing. But it is quite another thing to claim that at partial compliance or even *noncompliance*, I should not be required to do more than what others *would be* doing if they were in compliance. The first idea seems straightforward enough. The second idea is surely false: noncompliance *can* place added duties onto individuals.

Imagine, for example, four small children drowning in a shallow pool. In addition to yourself, there is only one other able-bodied swimmer.⁹⁰ Under the compliance condition, you are not required to rescue more than two children, since, *ceteris paribus*, demanding that you rescue more than two children would be demanding that you do more than your fair share. But now suppose that your counterpart simply refuses to discharge her responsibility: she can't be bothered to rescue any of the children. There can be no question that you are required to take up the slack and make a good faith effort to rescue those other two children. But according to the compliance condition, her noncompliance does not alter the scope of your responsibilities. You are not required to take up the slack left by her noncompliance. On Murphy's account, since full compliance is the benchmark against which your duties are fixed, you need only consider what your counterpart

⁸⁹ Myers (1999) argues that beneficence is at root a cooperative notion. This view is not convincing for reasons that I discuss in Chapter 8. This is not to say that cooperation has no place in the morality of rescue: indeed, as I argue in Chapter 7, it has a critical role to play.

⁹⁰ Assume, if you like, that neither I nor my counterpart bears any special relationship to the children.

would be required to do under full compliance in order to know what you are required to do under *any* circumstances. But surely this is wrong. Noncompliance can place added burdens at your feet.

Murphy concedes that there is a problem here. His response, however, is unpersuasive. He argues that even though my refusal to rescue the other two children would not be wrong, my motives would be open to rebuke. And, in any case, it is my appalling character—and not wrongdoing per se—that our intuitions are tracking in cases where I refuse to rescue the other children. Murphy (2000: 132) writes:

A person inclined to rescue the first [two children] would very likely also be strongly inclined to rescue the second [two]. In doing so, she may act beyond the call of duty, but she acts on motives she ought not try to rid herself of. And again, the person who did stop short with her fair share of the rescue load would stand out for having a highly undesirable character; this would be a case of ‘blameworthy right-doing.’

So while I may not violate any moral principles in refusing to rescue the other two children, I am nevertheless blameworthy for lacking beneficent motives. Setting aside the issue of blame, it’s difficult to swallow the claim that my refusal to rescue the other two children is an act of “right-doing.” At any rate, this effectively concedes the objection: noncompliance cannot alter our duties to rescue, despite what our considered moral judgments might say.

But why should an act of “right-doing” nevertheless be blameworthy? Why should you be blamed for acting rightly? The answer lies in the relation between the collective principle of beneficence and deliberation. Instead of appealing directly to the collective principle, agents should adopt (in Harean terms) an “intuitive level of moral thinking” that approximates the collective principle:

[T]he aim would be to hit on a pattern of motives and rules of thumb that will lead us to act roughly as well as we can at roughly the level of sacrifice that would bring our level of expected well-being down to what it would be under full compliance. And for the exceptional case, where our rules of thumb seem unreliable, we aim to act as well as we can in that particular case for roughly the level of sacrifice that would bring our level of well-being down to what it would be under full compliance from that point on (Murphy 2000: 120).

Because it is not feasible for agents like us to act directly from the collective principle of beneficence (at 178 words, the official version is, in Murphy’s words, a “monstrosity”), agents like us are expected to develop whatever motives and rules of thumb that will do the work of the collective principle. Hence, morality encourages the development of a general disposition to rescue those in need when doing so will not cost you more than it would under full compliance. But it is in virtue of the generality of this motive that we will not discriminate between two drowning children and four: from the standpoint of this general motive, the added level of sacrifice is

negligible. Therefore, while strictly speaking not wrong, your refusal to rescue the two other children indicates an absence of morally sanctioned motives.

This line of reasoning, however, is puzzling. If rules of thumb and motives earn in a place in moral cognition only because we are not “archangels” capable of calculating what our level of sacrifice would be under full compliance, then an individual (call him “Spock”) that could make such calculations could—indeed, *should*—dispense with the relevant rules of thumb and motives altogether, since he would be acting directly from the collective principle of beneficence. Significantly, for individuals like Spock there are no grounds for blaming (or praising) them for lacking the rules of thumb and motives that the rest of us must rely on. They can consult the collective principle directly. Indeed, it is arguable that Spock represents the moral ideal, according to Murphy’s collective principle, since his actions would perfectly accord with that principle. In any event, when Spock refuses to rescue the two other children, this is simply a case of “right-doing.” There is nothing in Murphy’s account that makes blame appropriate in this case. And this, I contend, makes Murphy’s account highly dubious.

But perhaps rules of thumb and motives are not dispensable in this way. Perhaps the point of a motive to rescue others in need when such need is conspicuous is to “override” considerations of compliance. Even if, strictly speaking, rescuing the two other children would constitute more than your fair share, an appropriate motive of beneficence ought to be sensitive to the significant value of making this added sacrifice. But now it is fair to ask: why should the deliverances of the collective principle of beneficence occasionally be overridden if that principle is supposed to

provide, according to Murphy, a complete and satisfactory account of our duties of beneficence? If morality recommends developing a motive that can override considerations of compliance, shouldn't this be captured in the principle itself?

The natural implication of the four-children rescue case is that noncompliance can alter our responsibilities to rescue. It is, therefore, misleading to suggest that the problem lies simply in our motives and not in the content of the principle itself. Moreover, there is a serious instability in proposing a norm for right action that agents (like us) should act on only through a motive which is specifically designed to override it.⁹¹ In sum, the compliance condition, while drawing important attention to the role of cooperation in beneficent undertakings, fails to account for the sense that partial (or non-) compliance can alter our reasons for providing rescue. A principle of beneficence that is insensitive to such cases ought to be rejected.

In this chapter we have considered two general attempts to limit our obligations to reduce suffering: arguments from the personal point of view and arguments from the impersonal point of view. Neither set of arguments could be made to work. The problem, as I will argue in the next chapter, lies upstream: what makes revisionism or any normative account that relies on an optimizing principle of beneficence implausible does not have to do with its failure to recognize limits on what we can be expected to do for others. Instead, the problem lies in the assumption that there is a standing reason to reduce suffering. After all, one of the background assumptions of this chapter was the assumption that in the absence of reasons to limit

⁹¹ This closely parallels discussions of indirect consequentialism. See Williams (1973); Scheffler (1983, Chapter 2); and Brink (1986).

our obligations to assist others, our standing reason to assist others becomes an overriding reason, which in effect amounts to a duty to assist others. (This follows from the schema introduced in Chapter 2.) But, as I argue in the next chapter, there is no standing reason to assist others, insofar as this is thought to be a reason whose content is specified by reference to the badness suffering or the goodness of happiness.

CHAPTER FIVE

Suffering and Practical Reason

TO REACH THE CONCLUSION that we have a conclusive reason to reduce as much suffering as possible, we must first assume that there is a standing reason to reduce suffering; we must assume that, at any given moment, there is a reason to reduce suffering.⁹² For only with such a reason in place does the failure to establish reasons that can, at least often, outweigh this reason render our standing reason conclusive. In Chapters 2 and 3, we explored the content of this reason, but did not question its existence. And, in Chapter 4, it was only because we assumed its existence that we sought ways of limiting its influence. The aim of the present chapter is to explore this standing reason to reduce suffering directly.

I will argue that there is no standing reason to reduce suffering. We can account for our considered moral judgments *without* having to posit a standing reason to reduce suffering. We need not accept the claim that, for example, there is a standing reason to promote the good, nor must we accept the claim that the badness of suffering is itself a reason to reduce (or prevent) it. At first blush, these denials sound untenable. But on closer inspection, it becomes clear that accepting them comes at a higher cost than rejecting them. The problem lies with the assumption that suffering is unconditionally bad, that its badness is not conditioned on anyone's having a reason for reducing it. In its place, I argue that suffering is bad, if it is, in virtue of the

⁹² Unless otherwise stated, "suffering" will be used to refer to a non-instrumental bad or disvalue, according to which, independently of whatever good might result from the suffering (e.g. retribution, inner strength), the experience of suffering is intrinsically bad, such that it would be better, either for the sufferer or simpliciter, that it did not occur.

badness *to* sufferers. Our reason to respond to the serious suffering of others has as its object sufferers, not suffering, for it is persons and their reasons that matter and not their experiences *per se*. Let me begin with a few general points about reasons.

If there is a standing reason to Φ —or, as it is sometimes put, a *pro tanto* reason to Φ —then there is a reason that always counts in favor of Φ -ing. Even when other countervailing reasons have the effect of outweighing one's reason to Φ in a particular setting, it remains the case that there is a reason to Φ . It's just that, in that setting, one's reason to Φ is not decisive. Standing reasons are not to be confused with *prima facie* reasons. Unlike the former, *prima facie* reasons can turn out not to be reasons at all.⁹³ For example, if it turns out that the liquid I am about to drink is in fact bleach (and not water as I had assumed), then what I took to be a reason to drink that liquid is in fact not a reason, since bleach is not likely to quench my thirst. Standing reasons do not cease to be reasons in this way. If there is a standing reason not to drive in dangerous conditions, then such a reason does not cease to be a reason even when I decide to drive in dangerous conditions because, for example, that's the only way I can get my seriously injured child to the hospital. In this case, the standing reason not to drive in dangerous conditions is outweighed by my reason to save my child's life. The point is that my reason not to drive in dangerous conditions does not cease to be reason even when I decide not to act on it. It continues to have, as it were, positive weight or positive force.

How, then, should we understand the claim that there is a standing reason to reduce suffering? First and foremost, it is a claim about when such a reason is a

⁹³ I'm aware that this gloss on *prima facie* reasons departs somewhat from standard usage. In thinking of *prima facie* reasons as *apparent* reasons, I follow Kagan (1989), p. 17.

reason for us: *always*. There is always a reason to reduce suffering. Formally speaking, for any action Φ , if Φ would reduce suffering either directly or indirectly, then we have a reason to Φ . Again, this is not the claim that the reason we have to reduce suffering is always decisive.⁹⁴ After all, the claim that there is a standing reason to reduce suffering is consistent with the claim that no one, as it happens, ever has a decisive reason to reduce suffering. For if the circumstances were such that countervailing reasons always override our standing reason to reduce suffering, then the reason to reduce suffering would never become decisive. The upshot, then, is that the claim that there is a standing reason to reduce suffering is not the revisionist claim that we almost always have a decisive reason to reduce suffering, even if revisionists typically rely on the former reason. Second, to say that we have a reason to reduce suffering (let alone a standing reason) is to say, roughly, that there is something good or worthwhile or valuable about reducing suffering.⁹⁵

Now there is, I grant, a considerable amount of controversy surrounding this last point. For example, some will want to say that having a reason to reduce suffering is just to say that reducing suffering would satisfy (perhaps after careful reflection) at least one of my desires.⁹⁶ It should be obvious, however, that on this reading it is very likely false that there is a standing reason to reduce suffering. It's implausible that for any agent at any given moment, the object of at least one of her desires is the reduction of suffering. But I do not intend to argue for this here. It is

⁹⁴ A quick note about my terminology: decisiveness, as I will be using the terms, refers to normative properties of reasons, not to properties of, say, human motivation. Specifically, a reason R is decisive just in case the existence of R is sufficient to make it the case that, all things considered, one ought to act on R.

⁹⁵ See Raz (1999) and Scanlon (1998), Chapter 1.

⁹⁶ For useful surveys of the range of positions, see Wallace (1999) and Pettit et. al (2004).

enough for our purposes to simply grant that having a reason to reduce suffering consists in the fact that there is something good or worthwhile or valuable about reducing suffering. For if there is any hope of defending the claim that there is a standing reason to reduce suffering, the claim will have to be understood along these lines, viz., that there is always something good or valuable or worthwhile about reducing suffering, independently of what anyone desires—in short, that there are always considerations that count in favor of reducing suffering.

I. A Standing Reason to Promote the Good

In Chapter 2 we explored the suggestion that our reason to rescue the drowning child rested on a more fundamental reason to promote the good. According to Kagan, our standing reason to reduce suffering is simply an extension of this more fundamental reason to promote the good. The thought is that, from the point of view of the universe, reducing (even minimally) the amount of suffering in the world has the effect of promoting the good insofar as the presence of suffering, *ceteris paribus*, decreases the good in the world. And Kagan (1989: 17), like consequentialists generally, “recognizes no other morally acceptable reasons at all,” apart from this reason to promote the good. We can think of the consequentialist then as proposing the following monistic hypothesis: there is always one and the same consideration that counts in favor of reducing suffering, viz., that, *ceteris paribus*, reducing suffering promotes the good.

I argued in Chapter 2 that the fact that rescuing the drowning child would promote the good does not adequately account for the content of our reason to rescue

him. As I suggested there, what's missing from the consequentialist's story is any account of how one's refusal to rescue the child represents a serious wrong *to the child*. On the consequentialist's approach, there is no sense in which you wrong the child directly. Instead, you wrong him, at most, indirectly—by reference to how his death diminishes the overall good. But surely there is such a sense. We fail to take his suffering directly into account when we refuse to rescue him. Kagan misdiagnoses the root of the problem. But the monistic hypothesis the consequentialist puts forth suffers from a more fatal flaw.

Recall the case of vintage boots from Chapter 2, where Peter is faced with the choice of rescuing the drowning child or selling his vintage boots for humanitarian ends. On the consequentialist hypothesis, Peter only has reason to do what will promote the overall good. Recall that apart from reasons to promote the overall good, the consequentialist recognizes “no other morally acceptable reasons at all.” So while it follows that Peter has a reason to sell the boots and refuse to rescue the child, it also follows that he has *no morally acceptable reason at all* to rescue the child, since, *ex hypothesi*, doing so would not promote the most overall good. But if Peter has reason to do anything at all, surely he has reason to rescue the child, even if it could be argued that, on balance, that reason is overridden by other reasons. But if this is right, then there are other reasons to rescue the child, reasons that do not appeal to the promotion of the overall good. And such reasons would entail the rejection of the monistic hypothesis put forth by the consequentialist, since if there is a standing reason to reduce suffering, it is sometimes supported by considerations that are independent of how much good can be promoted.

Kagan attempts to counter this objection. His reply involves a clarification: what is meant by the claim that we have a standing reason to promote the good is that we have a standing reason “to promote each individual good.” For each individual whose good is at stake, then, we have a corresponding reason to promote that good. But of course our “strongest reason will be generated in favor of the reaction that will promote the greatest amount of good overall.”⁹⁷ Therefore, the monistic hypothesis put forth by the consequentialist can support the idea that there is *some* reason to rescue the child since doing so would constitute the promotion of a good, even if it is not the only good at stake.

But this line of response is far from convincing. For Kagan’s regard for individual goods (mine, yours, his, hers) is always at the mercy of his regard for the overall good. One can—indeed, must—always dispense with the former in favor of the latter, but never the reverse. But this robs individuals of whatever distinctive worth they might possess.⁹⁸ According to the consequentialist’s picture, one’s worth is always a function of its contribution to the whole, but this is inconsistent both with our attitudes towards those with whom we interact and with our attitude towards our own worth.

As a rather loose analogy, consider a financier’s relationship to a single dollar. Her fortune is of course dependent on the accumulation of individual dollars, so, in that sense, every dollar does matter; hence, she has some reason to acquire this or that dollar. But what really matters is the *accumulation* of dollars, the lump sum. And this concern militates against the concern for any single dollar. After all, to make

⁹⁷ Kagan (1989), p. 57.

⁹⁸ The “distinctive worth” of individuals is sometimes captured by the notion of the “separateness of persons.” See Raz (1986), Chapter 11.

money, you have to spend it. So the suggestion that the financier's concern is with each dollar is highly misleading. Correlatively, the suggestion that our standing reason to promote the good consists in our reason to promote each individual good is highly misleading in the context of what would promote the greatest good overall. At the end of the day, it's populations—not persons—that count.

There is, at the same time, a more general mistake that Kagan falls prey to, viz., in thinking that the existence of a standing reason to promote the good is the “best explanation” of many of our considered moral judgments. For example, Kagan (rightly) assumes that we judge that we have a decisive reason to rescue the drowning child when doing so comes at little cost or risk to the agent. Kagan then has us imagine a scenario in which rescuing the child requires sacrificing our own life, such that we judge that there not a decisive reason to rescue the child. Here Kagan assumes (again, rightly) that we judge that there is nevertheless *some* reason to rescue the child, even though this reason is not decisive. We might explain this by saying that while the difference in cost can affect the requirement to rescue the child in the two cases, it does not affect “the fact that there is a reason—the same reason—to save the child in both cases.”⁹⁹ But then Kagan (wrongly) concludes that the best explanation for our judgments in these cases turns on the fact that there is a standing reason *to promote the good*. All of Kagan's examples are pressed into similar service: since our reason to reduce suffering does not disappear even when other considerations override this reason, there must be a standing reason to promote the good that best explains our judgments.¹⁰⁰

⁹⁹ Kagan (1989), p. 49.

¹⁰⁰ See *ibid.*, pp. 47-64.

But the only thing Kagan has shown, if he has shown anything, is that in these cases there is *some reason or other* that always counts in favor of reducing suffering. He has not shown that the content of this reason must be construed in terms of promoting the good. It does not occur to Kagan that there are other, equally viable, ways of characterizing a standing reason according to which we have a reason to reduce suffering. If there really were no other means of characterizing this standing reason, then Kagan would be right that we “cannot deny the pro tanto [i.e. standing] reason to promote the good without retreating into the minimalist’s camp,” according to which there is no reason at all to reduce suffering in these cases.¹⁰¹ But surely we can deny the standing reason to promote the good without having to admit that there is no reason at all to reduce suffering. We need only ascribe different content to that reason. For example, we can agree that even though one might be prohibited from killing one innocent person as a means of saving the lives of five others, there is nevertheless some reason to save those five lives. But we need not agree—nor should we agree, if I’m right—that the reason has to do with the promotion of the good. For all that’s been said, the reason may have to do with the inherent disvalue of losing one’s life or with respecting the rational agency of others or with responding in ways that reflect benevolence, to name just a few.

¹⁰¹ Ibid. p. 19. Discussing a later example involving thresholds, Kagan commits the same mistake: “If a skeptical moderate were to deny the existence of the pro tanto reason to promote the good, he would have to hold that there is no reason at all to save the ninety people, although there *would* (suddenly?) be a very powerful reason if the number were one hundred” (p. 51; italics in original). But again the consequent follows only on the assumption that the only way to construe this standing reason is in terms of promoting the good. I am denying this.

II. A Standing Reason to Reduce Suffering

There is, however, a more direct way of approaching our question, one which does not rely on the notion of promoting the good. To wit, there is always one and the same reason counting in favor reducing suffering, viz., that suffering itself is a bad and ought not to occur. Here, it is the direct appeal to the badness of suffering—and not its bearing on the promotion of the good—that grounds our reason to reduce or prevent it.

In Chapter 2, I criticized this view on the grounds that it left unanswered fundamental questions about the relationship between suffering and our reason to reduce it. As I argued there, anyone at all skeptical of the suggestion that suffering directly grounds a duty to abolish it will not be persuaded by an account like Mayerfeld's. For as Mayerfeld (1998) himself acknowledges, it is "not the sort of account for which [he] can supply an argument."¹⁰² Rather, his hunch is that we will find it obvious that "where there is suffering, there exists a demand or an appeal for the prevention of that suffering." And this demand issues directly from the badness of suffering: "suffering cries out for its own abolition or cancellation." But this suggestion, as evocative as it might be, fails to adequately justify the claim that we always have a reason to reduce suffering. The crucial claim Mayerfeld fails to establish is the following: "the fact that suffering ought not to occur and the prima facie duty to prevent suffering refer to the same moral phenomenon."¹⁰³ Any serious

¹⁰² All quotations in the present paragraph are taken from pp. 111-2.

¹⁰³ Mayerfeld uses the locution "prima facie duty" to refer to what I've been calling a standing reason (and what Kagan refers to as a pro tanto reason). That he has in mind standing reason is clear from such passages as: "When other moral considerations are present—other prima facie duties, or moral permissions not to incur certain costs—the prima facie duty to relieve suffering is sometimes *overridden* by them" (p. 113; italics added).

attempt to understand intentional action (by this I mean, acting for a reason), which would include reasons to respond to the suffering of others, must not take such claims at face value. For we are seeking an answer to a more fundamental question: what it is about another's suffering that provides us with a reason to end it? Mayerfeld leaves off precisely where we should begin.

Of course it would be hard to deny that the nature of suffering is such that, *ceteris paribus*, it would be better if it did not occur.¹⁰⁴ But better *for whom*, we might ask. Here, I would argue, Mayerfeld falls into the same trap that Kagan fell into: roughly, the concern for populations soon eclipses concern for persons.

According to Mayerfeld, suffering is not only bad for the person that experiences it, but it is also bad "from the impersonal perspective." This, Mayerfeld thinks, follows from the claim that suffering is intrinsically bad. But by assigning objective disvalue to suffering in this way, Mayerfeld detaches suffering from the individual that experiences it and makes room for the possibility of aggregating badness across persons. It is in this way that our reason to reduce suffering assumes, as Mayerfeld (1998: 119) puts it, a "cumulative character," by which he means that our reason to reduce suffering "builds up to a [prima facie] duty to achieve the *greatest possible* reduction in the cumulative badness of suffering. We do not comply fully with the duty to relieve suffering until we achieve maximum reduction."

But now Mayerfeld's position is open to the same objection leveled against Kagan: there is, at best, only a nominal reason to relieve the suffering of the drowning

¹⁰⁴ Among the sorts of cases the *ceteris paribus* clause is meant to rule out are cases where criminals are made to suffer for their crimes or where suffering is a necessary component of moral regeneration or "character-building." It might be argued that under these circumstances suffering indeed ought to occur. I will not pursue such questions here. See below for further discussion of the *ceteris paribus* clause.

child in the case of vintage boots. There is, we might concede, a poetic sense in which the child's suffering "calls out for its own abolition" and, in that sense, provides Peter with some reason to prevent it. But since Peter can prevent far more suffering by selling the boots and donating the proceeds to humanitarian relief efforts, and since "nothing less than full compliance is required of us," whatever reason Peter has to rescue the child is immediately silenced by his reasons to sell the boots. Hence, a monistic hypothesis that treats the badness of suffering as something that can be detached from persons in a way that allows for aggregation fails to account for other important considerations that bear in our relations to others in need.

Of course if morality had the structure that Kagan and Mayerfeld (and other consequentialists) envision, then it would come as no surprise that there is a standing reason to reduce suffering. But surely we can conceive of a morality that does not rest on a standing reason to reduce suffering. Such a morality might consist of a reason to treat others only in ways that respect their dignity, or that they could rationally consent to or that they could not reasonably reject. The point, however, is that reasons of this sort would ground a reason to render assistance in varying scenarios even though their contents make no direct reference to the reduction of suffering. Whether or not such an account can be made to work, there's no need to accept the claim that the only sort of morality worth countenancing is one that makes direct appeal to suffering or the diminution of the good. After all, non-consequentialists have always argued for alternative conceptions of morality, ones in which states of affairs are not the ultimate (or only) bearers of value, if they have value at all.

But perhaps we ought to step back from the suggestion that there is *always* a reason to reduce suffering. In the previous paragraph I suggested that there are ways of conceiving of morality that do not presuppose a *standing* reason to reduce suffering. Quite apart from this suggestion, however, what would it mean to say that there is always a reason to reduce suffering? Apparently, this would mean that for any given set of circumstances *C*, and for any individual *S*, *S* has a reason to reduce suffering in *C*. There is, in other words, always something good or worthwhile or valuable about reducing suffering in *C*, even if on balance there are other reasons overriding this reason. Now this suggestion seems credible when one of the available courses of action involves the direct reduction or prevention of suffering. But in what sense, if any, do I have a reason to reduce suffering when deciding between rice and beans for dinner, or deciding which shoe to put on first? Do I have a reason to reduce suffering at this very moment, in my attempt to articulate a particular philosophical thesis? Among the reasons that bear on my decision to have lunch now or work for another hour, should I include the fact that there is something good or worthwhile about the reduction of suffering? It's hard to know what to make of the suggestion. Many of our moment-to-moment deliberations appear to assign not just little positive force to such a reason, but no force whatsoever. From the standpoint of practical reason, such a reason is inoperable.

Now defenders of a standing reason to reduce suffering will be quick to point out that these examples, whatever else they may suggest, do not suggest that there is nothing good or valuable or worthwhile about reducing suffering. Even if in the midst of deciding on a dinner menu the question of suffering is remote, we do not

want to say that, for this reason, there is *no reason at all* to reduce suffering. Surely the thing to say is that had the circumstances been altered in such a way that deciding on a dinner menu did have a bearing on the amount of suffering in the world, then we would have a reason to act in such a way that that suffering was reduced. This point is sometimes captured in a *ceteris paribus* clause: when all else is equal, one has a reason to act in such a way as to reduce suffering. Not surprisingly, the standing reason entails the *ceteris paribus* claim.

But the introduction of a *ceteris paribus* clause hardly establishes the existence of a standing reason to reduce suffering. It is one thing to claim that there is always a reason to reduce suffering and quite another to claim that when the only factor that distinguishes two course of action concerns the reduction of suffering, one has a reason to choose whichever action will reduce suffering. The latter claim, unlike the former, may well follow from an account of morality that contains no standing reason to reduce suffering whatsoever. For example, it may be that the only way to treat other individuals as persons deserving of respect is to reduce their suffering when all is else equal; this is not because we have a standing reason to reduce suffering, but because we have a standing reason to treat others as persons deserving of respect. At any rate, that one's reasons to reduce suffering are co-extensive with one's reasons under *ceteris paribus* conditions leaves open what the source of those reasons are. What defenders of a standing reason to reduce suffering have not shown is that the best (not to mention the only) way of explaining our reasons to reduce suffering under *ceteris paribus* conditions is through appeal to a standing reason to reduce suffering.

III. Suffering and Sufferers

I want to return now to a theme that dominated my discussions of Kagan and Mayerfeld, viz., the separation of suffering from the individuals that experience it and the way in which this leads to aggregation. In order to complete my diagnosis, I need to say more about a) the relationship between suffering and the individuals that experience it and b) the bearing this has on our reasons to respond to others in need.

To prepare the way, let me summarize my conclusions: first, our reason to respond to the suffering of others is not based on the badness of suffering *per se*, but on the reasons persons have for wanting to avoid suffering and the fact that those reasons are reasons we can share. That this should be the case is explained by the fact that, second, value supervenes on the relationship between persons and their reasons for pursuing or avoiding certain states of affairs, not on states of affairs *per se*.

The claim that value is a property supervening on the relationship between an agent and her reasons for pursuing or avoiding certain states of affairs needs to be approached carefully. The first order of business is showing that this account of value is indeed distinct from the one offered by consequentialists like Kagan and Mayerfeld. This is the task I will set myself to in the remaining part of this chapter. Part of this task consists in showing that this account is not just coherent, but can successfully account for many of our considered judgments about value. Moreover, this account is less ontologically contentious than its rivals, and so considerations of parsimony favor this approach.

It is important, however, not to lose sight of our aim in all of this, viz., to clarify the nature of our reason(s) to respond to the suffering of others. And we should not lose sight of the fact that the approaches we have been considering, notwithstanding my objections, are not entirely implausible in suggesting that these reasons consist in the fact that suffering is bad in itself (Mayerfeld) or that the diminution of the good is bad in itself (Kagan). Indeed, common sense may seem to make no room for another approach. I believe, however, that there is room for another approach—one that makes better sense both at the theoretical level, and, difficult as it might seem at this stage, at the intuitive level. But to make this case requires getting clearer on the nature of value itself.

Let us look more closely at the suggestion that our reason to reduce suffering is rooted in the idea that suffering is bad in itself. This idea can be expressed in a number of ways. For example, it is sometimes asserted that suffering is unconditionally bad, where this means that the badness of suffering is not conditioned on anything else's being bad. The badness of suffering is constitutive of the experience of suffering itself. A related but distinct suggestion is that suffering has agent-neutral or impersonal disvalue; the idea here is that anyone has reason to want to avoid or end suffering, and this would include "someone considering the world in detachment from the perspective of any particular person."¹⁰⁵ So if suffering has agent-neutral disvalue, then your reason to reduce suffering is not relative to what you (or anyone) happens to value. (It is little wonder then why some see the move from the intrinsic badness of suffering to a standing reason to reduce it as so compelling:

¹⁰⁵ Nagel (1986), p. 153. See Parfit (1984), p. 143, for the original distinction.

on this interpretation they come to the same thing.¹⁰⁶) At any rate, the point to bear in mind is that the intrinsic disvalue of suffering is supposed to be independent of either anyone's interest in reducing it.

This construal of the disvalue of suffering opens up the following line of thought: since suffering is impersonal, in the sense that its disvalue does not rest on *whose* suffering it is, there is no in principle obstacle to summing the badness of suffering across persons.¹⁰⁷ That is, by detaching suffering from persons, by treating their suffering as a property whose value can be compared objectively, we can sum your suffering and my suffering since they are, formally speaking, comparable. Again, this is made possible by the claim that, as Nagel (1986: 161) puts it, "suffering is a bad thing, period, and not just for the sufferer." Consider how these considerations are reflected in our intuitive judgments.

You suddenly become aware of a toddler drowning in East Lake and, at the same time, three toddlers drowning in West Lake. As it happens, you have only enough time to reach one of the lakes and no other help is available. Suppose as well that you bear no special relations to any of the children. Our intuitive judgment is that you should go to West Lake, where you stand a good chance of saving three children. As justification for this judgment, one might reasonably claim that more suffering would be prevented by saving the three children than the one. This is not to deny that the one child's death would be a tragedy. It is, instead, to assert that *together* the three deaths represent a greater amount of suffering than the one death,

¹⁰⁶ See Mayerfeld (1998), p. 112.

¹⁰⁷ If cardinal rankings are not warranted, perhaps it is enough that ordinal rankings are. Since the badness of suffering is independent of the identity of the sufferer, then your suffering and my suffering, being two tokens of the same type, represent a greater amount of suffering than the amount of suffering involved if it were only your suffering.

and it is on this basis that one has a conclusive reason to go to West Lake. But such an assertion is tantamount to claiming that the individual suffering of each of the children can be summed in such a way as to provide a means of comparing the relative amounts of suffering the two courses of action are likely to yield.

But recall that it was precisely this move to aggregation that led to the wrong response in the case of vintage boots, viz., that there is no compelling reason to rescue the drowning child. For the reason to save the drowning child, if there is one, is immediately and conclusively overridden by the reason to refuse his rescue in favor of the *collective* reduction of suffering of individuals abroad. But surely you wrong that child by refusing to save him even if it comes at a cost to the overall good. The vintage boots case suggests that there are occasions in which an individual's suffering should not be detached from him, in which one's plight should be viewed as incommensurate with the plight of others. In the case of vintage boots, your relationship to that child makes his suffering of special concern to you.

What is needed then is an account that has the consequence of rejecting the idea that the badness of suffering can be summed across persons. But to do this requires rejecting a more basic idea, viz., that suffering is intrinsically disvaluable, that the disvalue of suffering is not conditioned on anyone recognizing or caring about the reduction of it. *Prima facie*, this idea may not inspire much confidence; if anything is unconditionally bad, surely it must be human suffering. But this idea faces serious problems, some of which I hope to lay out below. In its place, I want to argue that suffering is bad, if it is, because it is a bad *for* the sufferer, because the

sufferer cares about its removal.¹⁰⁸ In this respect, suffering cannot be detached from persons. It is a mistake then to think that suffering has, as Nagel (1986: 160) claims, “a life of its own” and, hence, “a value of its own.” Since suffering does not have a disvalue independently of the sufferers that experience it, there is no straightforward way of summing the badness of suffering across persons.

As these remarks imply, the ultimate source of value lies in persons, not in their ends. I hasten to add, however, that this does *not* mean that we have no reason to promote the ends of others, e.g., the removal or reduction of their suffering. Rather, we have a reason to promote the ends of others, if we do, because this is one way of respecting the value of persons. Ends matter, if they do, because persons do.¹⁰⁹ Ends are extrinsically valuable, in the sense that the *source* of their value comes from persons who are intrinsically or unconditionally valuable.¹¹⁰ This emphasis on persons requires careful unpacking. For it might be wondered, if persons are intrinsically valuable, why not allow for the possibility of aggregating this value? After all, by doing so we could easily account for our judgment in the two lakes case: our reason to go to West Lake (and save the three) overrides our reason to go to East Lake (and save the one). In short, in what sense does this alternative differ from the consequentialist approaches we’ve been considering? For surely it is a misrepresentation of consequentialism to say that persons don’t matter or that persons

¹⁰⁸ My account of value shares important features with Korsgaard (1993). Korsgaard, however, does little to develop the idea of “shared reasons” which I hope to do here. More importantly, I am less sanguine than Korsgaard about the range of reasons we do or can share.

¹⁰⁹ Nagel’s (1986) account is actually ambiguous on this front: on some occasions he speaks of agent-neutral value, which suggests that the ultimate bearer of value are states of affairs. On other occasions, value is analyzed in terms of agent-neutral reasons, which would be closer to the account I favor. See Korsgaard (1993).

¹¹⁰ Extrinsic goods should not be confused with instrumental goods; they can be objectively good so long as the conditions for their goodness are met. See Korsgaard (1983) for a stimulating discussion of these distinctions.

matter only because their ends matter. As already noted, Kagan (1989: 58) is quick to counter this suggestion: “the overall good simply *is* the overall well-being of individual persons.”

But this commitment to the value of persons is in tension with a more basic commitment to impartiality, and it is this latter commitment that distinguishes consequentialism most sharply from the alternative I favor. Strict impartiality, for example, is behind the push to adopt the *impersonal point of view*, from which we speak of better or worse states of affairs, where the identities of affected individuals are ignored. The extent to which this militates against the value of persons has already been discussed at length. There is, however, another puzzle that arises in the context of the impersonal point of view.

Consequentialists are committed to the claim that there is a source of value that is not derived from the value *to* particular agents. If value does not have this character, then we could not claim that, for example, suffering is bad in itself and not just bad for the sufferer. We must assume then that there is a perspective from which suffering matters, but does not matter *to* anyone. Instead, it matters “from the perspective of the universe.” But this idea should strike us as deeply puzzling. For from the impersonal point of view, things matter not just because they matter to me and to you and to any other particular agent whose interests are at risk, but because they matter *simpliciter*.¹¹¹ Because this claim is critical to the coherency of the consequentialist program, it’s important to get straight on just what this idea implies.

¹¹¹ Koorsgaard (1993) uses the term “absolutely” to describe this view. Perhaps a better term would be “objectively,” where the implication (that I want to criticize) is that things can matter even if there is no one to whom they matter.

From the impersonal point of view, there is supposed to be a source of value that is over and above the value *to* anyone.¹¹² When, for example, my suffering is abated, what matters from the impersonal point of view is only that *someone's* suffering is abated, it matters not whose. The reason is that the alleviation of suffering is good *simpliciter* and not just good for me. Alternatively, that it is A's suffering that is abated and not B's is, *ceteris paribus*, irrelevant from the impersonal point of view, since what matters is only that *someone's* suffering is abated. The reason is that suffering is supposed to be bad without it necessarily being bad *for* anyone at all. But this is hard to make sense of. For surely it matters to A and it matters to B whose suffering is abated, so there is no confusion here about where (at least) one source of value lies. But what sort of judgment is it that says that independently of how it matters to A and to B, it *also* matters that someone's suffering is abated? It's not clear what the basis for this claim is, other than a reiteration of the consequentialist doctrine.

Why, in short, should we accept the claim that things can matter without regard to their mattering *to* anyone? After all, it seems but a short step from accepting this idea to accepting a form of Platonism, according to which being valuable is a property whose existence is independent not only of any particular mind, but of all minds. This sort of ontological independence, it seems to me, should be resisted.¹¹³

¹¹² See Foot (1988). Much of the inspiration for this line of attack can be traced back to Foot's remarks in that important piece. See also Foot (2001), Chapter 4.

¹¹³ See J. L. Mackie (1977). Although Nagel, apparently, wants to deny this, he is not always careful to block this implication: see his (1986), Chapter 8.

Of course if the judgment that it matters that someone's suffering is abated is the judgment of a consequentialist, then we might infer that it matters *to the consequentialist*—that is, to someone who sees value in the improvement of the whole—that someone's suffering is abated. But surely this should not be confused with the impersonal point of view. For even the consequentialist's point of view is not privileged from the impersonal point of view. More importantly, this answer will not satisfy anyone who desires to know how things can matter without regard to individuals for whom they matter.

In sum, the move to the impersonal point of view is, as it were, a move away from persons as such to ends. But this, I contend, is a move in the wrong direction. The basic source of value is not pleasure as such or desire-satisfaction as such, but rather how persons regard pleasure or pain or, for that matter, desire-satisfaction or whatever. Ends, if they are valuable, are valuable because they're valuable *for* individuals. Without reference to persons and the reasons persons have for pursuing a given project, we cannot say what the value of a particular project or relationship is. If this is right, then the impersonal point of view is not a perspective from which to appreciate value because that perspective essentially involves abstracting away from persons and the reasons they have for pursuing or avoiding various states of affairs, leaving only various states of affairs whose significance is left dangling. A more perspicacious account, then, would locate the source of value in the relationship between persons and their reasons for pursuing various ends. But this is not all. Others must be able to share these reasons, even if they do not share the ends. This is why value is said to be *intersubjective* and not merely subjective, according to which

whatever I happen to value *ipso facto* has value.¹¹⁴ This brief sketch raises a host of questions, some of which I should address here.

IV. Initial Objections

It might be wondered, for example, what difference, if any, these differences make in application. In other words, does the non-objectivist approach give us a reason to respond to others in need that the consequentialist doesn't? Suppose, for example, that I am right; suppose that our reason to respond to others in need is, at base, a reason to respond to persons as such. It is out of respect for persons that we have a reason to respond to their needs. But since all who experience suffering have reason to want it abated, and since the consequentialist maintains that suffering is itself a bad that ought to be abated, what difference does it make to speak of responding to sufferers rather than, say, suffering? If we assume, plausibly enough, that the set of individuals who have reason to want their suffering abated is co-extensional with the set of individuals that suffer, then what advantage is achieved by appealing to sufferers' *reasons*? It is—is it not?—the value of abating suffering that ultimately concerns us, not responding to others' reasons.

For example, when you confront a child drowning in your presence and you can save his life without substantial cost or risk, surely your reason to rescue him appeals, first and foremost, to the fact that *he will lose his life without your help*. We can safely assume that one of the child's ends is to avoid pain and suffering, even if it is not an end he himself could express. His drowning then represents a serious threat to one of his ends. Now the consequentialist maintains that it is this serious threat

¹¹⁴ Korsgaard (1993), pp. 28 ff.

itself that provides you with a reason to intervene. I maintain, on the other hand, that the serious threat does not itself provide us with a reason to rescue him; rather, it is the fact that the child does or should have reason to want to avoid the state that he's in, and this is a reason we can all share. In short, our reason does not track the threat to the child's life *per se*, but the child's *attitude* towards this threat. But surely this is not a plausible accounting of our reasoning.

For some of us, suffering, when we are unfortunate enough to confront it directly, tends to empty the mind of all thoughts unrelated to the suffering. We are left thinking: it should stop. The thought that the sufferer has an interest in having the suffering removed, while undoubtedly true, is distracting at best, and irrelevant at worst. After all, a person's interest in avoiding suffering is not just another interest someone happens to possess. It is not just one end among others. Avoiding suffering must surely be one of the most basic interests a person possesses.¹¹⁵ But it is a mistake to conclude from this that it is *this* interest in avoiding suffering that grounds our reason to end the suffering.

It's not clear, however, that this point constitutes an objection to my account. After all, we can admit that there is and always will be variation in the sorts of reasons people recognize to respond to the suffering of others. For that matter, there will be those that respond without the need to consider the reasons for responding at all. And we can admit that such a disposition deserves admiration. The question we are pursuing, however, is not a question of psychology, of what sorts of reasons

¹¹⁵ It is important to emphasize again that the sort of disvalue at issue here is non-instrumental disvalue. Most of us are aware of the instrumental roles suffering plays in, say, Christian doctrines of redemption or character-building, more generally. But this should take away from the point that avoiding suffering *for the sake of avoiding suffering* is a basic interest of all persons.

people actually recognize as counting in favor of responding to others' needs. The question we are pursuing is the normative one: what value is there in the reduction of suffering? That is, what sorts of considerations count in favor of reducing suffering? And a satisfactory answer to this question does not depend on people *actually* taking such considerations into account. This does not mean that we should ignore how people actually reason in this context. Rather, we can take stock of the sorts of considerations, if any, that lead people to respond to suffering in the way they do *without* having to admit that those considerations do, as a matter of normative fact, count in favor of the reduction of suffering. This is, at base, an inquiry into the nature of value. So I need not be troubled by the contention that an agent's response to suffering does not rest on considerations having to do with the sufferer and her reasons for seeking an end to her suffering.

But apart from the implications of this contention, can we be certain that the contention itself is well grounded? Is it obvious that our reasons to respond are in fact independent of the attitudes of the sufferers themselves, that is, independent of the extent to which the suffering is a bad *to* them, if indeed it is? One reason to think the answer to these questions is 'no' follows from considering cases in which we have a reason *not* to end someone's suffering or, outrageous as it may sound, to promote someone's suffering.

One advantage of my approach is that it can explain how such cases are possible. Recall that the objectivist about suffering (as she might be called) asserts that the badness of suffering is independent of anyone's reasons for ending or preventing it. But if the badness of suffering is independent of anyone's reasons for

ending or preventing it, then, presumably, the badness is also independent of anyone's reason for *not* ending it or even promoting it. But it is not hard to find cases in which suffering should not to be prevented, and not necessarily as a means of preventing future suffering. Some suffering should not to be prevented precisely because someone has good reason for experiencing it. If we are prepared to accept such cases, then we need to rethink the thesis that the badness of suffering is wholly independent of the relation sufferers stand to their suffering.

Consider, for example, cases in which one seeks to know the truth about a certain event, fully aware that knowing the truth could exact a very painful cost. For example, a wife desires to know whether or not her husband, around whom she has built an entire life, is involved in an extramarital affair. Imagine that you know that he is, and you possess evidence to convince her of this. But as it happens, only you are aware of the infidelity; as such, it's not unlikely that if she does not learn the truth through you, she may never learn it. Now it should be emphasized that this revelation would be devastating to her. By all accounts, she would suffer, and her suffering would be deep and lasting. So what should you do when confronted with the question, 'is my husband having an affair'?

On the face of it, a defender of the view that we ought to prevent suffering since suffering is bad in itself would recommend lying: 'No, he's not' or 'I have no idea.' If one has good reason to think that the fact of the infidelity may never surface again, then this would surely be the right course of action. But are we so sure?

Quite apart from whatever reasons you might have to avoid deceiving others, you also have a reason to tell her the truth, even though you can be sure she will

suffer dearly upon learning it: viz., she has reasons for wanting to know. That she and you both know that she will suffer, while surely relevant, is not what's at stake from her perspective. From her perspective, it is the truth that matters. And it is out of your respect for her and her reasons for wanting to know the truth that you should not withhold the information about her husband's affair. Even if we suppose that she is deceived about the mental anguish that lies ahead, and we suppose, further, that she will come to regret learning of his affair and, by extension, your revealing it, your reason for revealing what you know is nevertheless grounded in respect for her as an autonomous agent.

But the reference to 'respect' ought not mislead us. The point is that in weighing the reasons that bear on your actions, it is surely appropriate to consider what *she* has reason to want. And the force of this consideration is independent of its bearing on her suffering. That is, you do not consider what she has reason to want simply because in so doing you will be lessening her suffering. You consider what she has reason to want because such reasons are for you (and anyone else under the circumstances) normatively significant.

Let's probe a bit deeper. It might be objected, for example, that your reason for revealing what you know is in fact grounded on the expectation that in so doing you will decrease the likelihood of greater future suffering. There are several ways this thought might be cashed out. One might argue, for example, that the longer she lives under the illusion of her husband's fidelity the more painful the discovery will later be. Even supposing that the probability of that discovery is small, we can assume that the suffering associated with that discovery is of such a size that it would

be wiser to avoid the possibility of a calamity in the future by acting in the present. Another approach might be to argue that one's well-being is determined in part by the degree to which one is successful in one's projects and relationships.¹¹⁶ Since, *ex hypothesi*, the wife is being deceived by her husband, and since this threatens the success of one her most important pursuits, you would be contributing to a decrease in her well-being by not revealing what you know. Though it may not be in her self-interest to discover such a fact, her life as a whole would be worse off. In any event, on this approach, it is not by reference to her reasons for wanting to know the truth that grounds your reason to reveal what you know.

Finally, it might be argued that your reason to reveal the truth rests on your reason to see it to it that her desires (or her rational desires) are satisfied. Because, on this approach, her well-being depends on the extent to which her desires are satisfied, and since one of her strongly held desires is not to be deceived, you have reason to reveal what you know. While not directly dependent on suffering, this response nevertheless takes features of well-being as basic to our reasoning.

None of these suggestions, however, seems especially plausible, particularly because each of them seems to miss the essential point of the example, viz., that you are being asked to respond *to her*—not to the amount of suffering she is likely to experience, not to the fluctuations in her well-being, not even to the number of desires that you happen to satisfy. I do not want to deny that each of these proposals gets something right, nor do I want to deny that such considerations may never enter into our thinking. After all, it would be foolish to raise doubts about the extent to which

¹¹⁶ See Raz (1986), Chapter 12. Though Raz has endorsed this account of well-being, he would not endorse the line of reasoning in which his account is embedded.

she wishes to avoid suffering or avoid being deceived. However, if one carefully inspects the content of each of these reasons, one finds that there is something missing—or, more to the point, that someone is missing. When the wife asks you for the truth about her husband, she is asking you to do something *for her*, not for her well-being or for her preference ranking, as important as these things are. For she is distinct from these things; *she is the person to whom such things matter*. And that should be the object of your reason to tell the truth in those circumstances.

Now even if these remarks are on the right track, the objectivist can reasonably respond by pointing out that her potential suffering, however she views it, gives you *some* reason to prevent it, even if in the end your reasons to allow such suffering override that reason. What my remarks suggest, argues my critic, is not that suffering is not bad in itself, that suffering is not intrinsically disvaluable. What my remarks suggest is merely that in addition to the intrinsic badness of suffering there are considerations that do not themselves make appeal to suffering, and which may have important bearing on the reasons one has to reveal the truth. In support of this idea, we might alter our example in the following way.

Instead of the revelation leading to disappointment, suppose that it were to lead to death. You know, for example, that the husband, fearing financial ruin and public humiliation, would rather risk having his wife murdered than allowing her to live with that knowledge, since (we'll suppose) he fears she may use it against him. Now if such a threat was credible, it seems plausible that you would have conclusive reason to withhold your knowledge from the wife. But if one's reasons are supposed to track the person and her reasons and not states of affairs (e.g. her suffering or

potential suffering), then this case would appear to represent a counterexample to that suggestion. For it would appear that one's conclusive reason to withhold your knowledge in this case does refer to her potential suffering. That she has reason to want to know the truth, while relevant, is overridden in this case. Moreover, if the force of your reason to prevent her suffering can override other reasons in a case like this, then surely it is plausible to think that it retains some force even in those cases in which it happens to be overridden.

But this example does not establish what it sets out to establish. After all, nothing prevents us from appealing to the reasons persons *would have* were they in possession of the relevant information. If the wife were she apprised of the fact that by coming to know that *p* she would be putting her life at risk, then she would have a compelling reason to want to avoid coming to know that *p*. In the situation at hand of course, you cannot guarantee that she is in possession of such a reason. Indeed, the possibility exists that her reasons for wanting to know the truth are such that she would be willing to run the risk of losing her life. But this only reinforces the strength of my case, since if you were in possession of such an "advanced directive," you would have a compelling reason to reveal what you know. At any rate, our deliberations reveal that a central part of our thinking is reserved for considerations having to do with the reason agents have for wanting to pursue or avoid various states of affairs. And this is precisely the thrust of my position. As to the suggestion that there is some reason to prevent her suffering, even if it is not a decisive reason, we will recognize this as the suggestion that there is a standing reason to prevent or

reduce suffering, an idea we have already considered and for which we found little support.

It must be admitted, however, that the idea that suffering itself is a bad that ought to be prevented is not easily dispelled. The discrimination between sufferers and suffering, while perhaps tenable, does not come naturally, and so the suspicion lingers that the intersubjectivist has missed something. Of course we should not be put off by the possibility that the objective disvalue of suffering is an illusion whose grips on us remains even after the illusion is exposed. It may be that other facts about reason simply dispose us to think in those terms. As far as it goes, there may be no conclusive reason to counter this disposition.

In this chapter, I've argued against the idea of a standing reason to reduce or prevent suffering. This claim should be kept separate from the claim that we have no reason to reduce or prevent suffering, a claim I reject. What I have argued is that our reasons to reduce or prevent suffering do not all flow from a single concern, like the badness of suffering. Part of the problem lies in divorcing suffering from the sufferer and suggesting that it is the suffering to which we ought to be responding. Instead, duties of beneficence, assuming there are such duties, are sensitive to the reasons individuals have or would have under the circumstances. There may of course be limitations on what we can know about others' reasons, but this is not an objection to the idea that, ultimately, what matters are persons, and not their experiences, even if their experiences are a critical part of how persons see themselves. But this idea has implications beyond the idea of suffering.

CHAPTER SIX

Unique Dependency

HAVING ARGUED AGAINST THE IDEA of a standing reason to reduce suffering, my task now is to describe the conditions under which we do have a duty—a perfect duty—to rescue others in need. In this chapter, I offer the primary condition for rescue, one that we looked at in Chapters 2 and 3: unique dependency. If someone is uniquely dependent on me for rescue, if, that is, I am the only person in a position to provide the necessary assistance in an emergency situation and providing such assistance does not come at significant cost or risk to myself, then I have a moral obligation to make a reasonable effort to provide such assistance. Refusing to do so is wrong. The aim of the present chapter is to provide the moral foundations for this view, in particular, the foundations for the idea that I have a greater responsibility for those who are uniquely dependent on me than those who are not so dependent on me. As I will argue, my moral relationship to someone uniquely dependent on me is characterized by the expectation that I will not see the value of assisting him as interchangeable with the value of assisting others. There is a second-personal aspect to unique dependency that sets it apart from humanitarianism in general.

This approach to rescue brings together several ideas that we have already encountered. First, there is the issue of determinacy. Unique dependency entails determinacy. The significance of determinacy emerged in Chapter 3, where I attempted to show that one has a greater moral responsibility for determinate individuals than indeterminate individuals. In the next chapter, a more thorough

accounting of moral justification will show why determinacy deserves the attention it has been given. Roughly speaking, moral principles can only be accepted and rejected by individuals. Groups are prohibited from aggregating their individual complaints against a proposed principle in such a way as to swamp the significant complaint an individual may have to an opposing principle. Determinacy then will make an important difference in assessing the justifiability of proposed moral principles.

Second, as emphasized in previous chapters, to capture the moral significance of rescue it is not enough to point to the harm that I can easily prevent.¹¹⁷ While the harm is no doubt unfortunate for the individual at risk, the moral force of my reason to respond lies in my *relationship* to the imperiled individual: that the circumstances have selected me and only me for action implies that the imperiled individual has a claim on me that he can have on no one else, such that any reason I might have for not making the effort to rescue him must meet a very high standard, one that *he* could not reasonably reject. Thus when I refuse to take the minimal steps necessary to prevent the harm, our moral condemnation makes essential reference to the way in which I was related to the victim, to the fact that I was in control of his situation. But this is not simply a way of linking me with the badness I did nothing to prevent. Rather, moral condemnation links me with the *individual*—not what was done to the individual. This reinforces an idea touched on in the previous chapter.

Third, I want to insist that our duties to respond to distant suffering, if any, are fundamentally distinct from our duties to respond to emergencies. And the distinction lies in our different relationships to potential individuals. One of the

¹¹⁷ Although we expect that, as practical matter, that should be enough to spur you into action.

crucial differences between refusing to respond to individuals uniquely dependent on us and refusing to contribute to humanitarian relief efforts has to do with what individuals can be responsible for: in the former case, responsibility follows from the fact that had I taken the minimal steps required to save his life, *that individual* would almost assuredly be alive today. By contrast, in the latter case, had I made a contribution to humanitarian relief, there is no determinate individual we could point to whose life would have been saved through my contribution. What I am responsible for in the case of humanitarian relief is not making an effort to reduce suffering abroad, but this is to be distinguished from responsibility for any one (determinate) individual's death. To be sure, this is not the claim that my contribution to humanitarian relief does no good. It is that the good it does cannot be measured in (determinate) lives saved.

Finally, as should already be clear, a large part of the motivation for my account comes from careful consideration of what failing to respond to emergencies implies. As I argued in Chapters 2 and 4, to understand the kind of violation you commit when you refuse to provide assistance, we cannot simply refer to the harm that you could have been prevented, for this leaves out any essential reference to the *individual*, that is, the individual who experiences the harm. I have argued that suffering is bad, if it is, because it is bad *for* a person; it is not bad in itself, as Nagel and others argue. The difference is apparent in the **Vintage Boots** (Chapter 2: p. 29), where Peter decides to ignore the drowning child before him so that he can sell the boots for humanitarian relief abroad. If harm were the only relevant factor in deciding what to do, then Peter does the right thing in refusing to rescue the drowning child. But most of us find Peter's actions appalling, even if our reasons are not

immediately obvious. Therefore, since harm is not the only relevant factor, we must look elsewhere for the source of our judgment. I have argued that once we distinguish suffering from the individual that suffers, it is possible to see where Peter errs: he fails to see that his relationship to the child, under the circumstances, is not interchangeable with preventing harm elsewhere, since his primary obligation flows from his obligations to *this* individual, to the child, not to preventing harm generally. What needs explaining then is the idea that being in a unique position to provide assistance places significant moral constraints on conduct. In the next section I offer just such an account. In section II, I connect these remarks to a more fundamental idea, viz., that morality itself concerns our how we relate to other persons, and in the final section I deal with specific objections to my proposal.

I. Incommensurability, Dignity, and Dependency

What troubles many people about **Vintage Boots** is the idea that Peter would even *contemplate* the relative values of rescuing the child versus selling the boots. There are at least two ways this concern might be spelled out. On the one hand, the concern might be directed at the way Peter himself reasons under the circumstances. That he would calculate the relative values of the various outcomes, it might be argued, reflects a troubling disposition in Peter—one that act utilitarians, for example, have long wrestled with.¹¹⁸ On the other hand, the concern might be directed at the more general thought that those relative values should not be compared, not because it reflects a vicious disposition, but because it is inconsistent with the proper

¹¹⁸ See Brink (1986) and Railton (1984).

appreciation of persons.¹¹⁹ It is this second concern that I wish to explore, for I believe it goes to the heart of the matter.

As I will argue, the importance we place on unique dependency in our moral deliberations reflects in part our belief in the dignity of persons and, by extension, our relationship to other persons: appreciating the dignity of persons requires a general and abiding reluctance¹²⁰ to make the sorts of value comparisons exemplified in **Vintage Boots**. If it is possible to deliver a plausible, coherent picture of the significance of unique dependency, then we should take seriously the idea that there is a distinct category of moral reasons for rescue, one that accounts for the intuitive distinction between good Samaritans and humanitarians.

In the first stage of my argument, I consider some of Joseph Raz's remarks concerning incommensurability. According to Raz, in order to appreciate the significance of certain relationships (he mentions parenthood and companionship), one must judge that the value of that relationship is incommensurable with the value of other significant pursuits. Furthermore, one must be loathe even to consider exchanging one incommensurate option with another. In the second stage, I connect these remarks with Kant's discussion of the value of humanity in the Groundwork, in particular, the notion that humanity, as an end in itself, possesses dignity. There is, it seems to me, an important thought that unites Raz and Kant: We cannot appreciate the value of other persons and our relationships to them if we are generally prepared

¹¹⁹ The difference is a difference in emphasis: the first interpretation emphasizes aretaic features of Peter, whereas the second emphasizes how one is to conduct oneself in response to different sources of value.

¹²⁰ That our reluctance is "general and abiding" is meant to reflect the fact it is, after all, *reluctance* and not *refusal* that I'm advocating. For it is one thing for Sophie to be reluctant to consider which child she ought to save; it is another for Sophie to refuse to consider the choice at all. See below.

to consider comparing (let alone, exchanging) that value with the value of other pursuits. And this thought provides a clue to understanding the significance of the relation of unique dependence and, by extension, of our relations to distant others.

Recall that the relation of unique dependence is marked by one person's reliance on one *and only one* other person. If Shelley is uniquely dependent on Peter, then Shelley's continued survival depends critically on Peter's and only Peter's assistance. If Peter refuses to assist Shelley, Shelley almost certainly dies. Peter thus has a choice that Shelley does not, a choice that unambiguously determines the survival of another person. Now when those choices are a) rescue Shelley (at little cost or risk to Peter) or b) refuse to rescue Shelley merely for the sake of convenience, Peter should choose (a). But deciding between these choices does not require much in the way of moral discrimination. After all, it is a matter of moral common sense that the right thing to do under the circumstances would be to rescue Shelley.

But there is a deeper point here, one that concerns practical thinking more generally and which shapes our focus in this section: A morally sensitive agent is reluctant even to deliberate about the sorts of considerations that count in favor of *not* rescuing Shelley under the circumstances. No doubt Peter has the choice of not rescuing Shelley, but in conceding this we should not be understood as implying that such a choice is really "on the table." For most of us, the choice is immediately

recognized as one we could not “live with,” to use Williams’ (1973) phrase.¹²¹ We cannot, as such, identify with this option.

Analogously, it is misleading to say, in reference to me: “He does not think he has reason to murder his rich uncle.” For this implies that, independently of the query, I had reason even to examine my reasons for murdering my rich uncle. What calls out for attention here is not my reason for not murdering my rich uncle, but rather *the reason for posing the question in the first place*. Thus the point is not simply that Peter is not expected to adduce reasons for not choosing the second of his two options, i.e., not rescuing Shelley. The point is that Peter should be loathe even to deliberate about the considerations, if any, that count in favor of not rescuing Shelley.

But what happens when we introduce a further choice, viz., c) refuse to rescue the child in order to sell his boots for humanitarian ends? Now it will be observed, if not emphasized, that there is *some* consideration that counts in favor of Peter’s selling the boots, viz., relieving distant suffering. After all, just because Peter’s plans are interrupted by the drowning child does not mean that his reason to sell the boots suddenly ceases to be a reason; it is still a reason. It is just that, under the circumstances, that reason is overridden by other reasons.

This response, however, falls afoul of the same confusion above: While it is true enough that Peter’s reason to sell the boots for humanitarian ends is overridden by other considerations, this is not the relevant issue. The relevant issue concerns our general mistrust of a certain mode of practical thinking, viz., Peter’s willingness to

¹²¹ My use of the notion of “live with” departs from Williams’ usage insofar as it is not one’s integrity that is at issue so much as one’s *moral self-respect*, the degree to which one could reconcile one’s actions with one’s conception of oneself as a decent person.

treat the child's plight as one disvalue among others, given that the child is uniquely dependent on Peter. But what, if anything, does this mistrust indicate?

Incommensurability

Raz argues for the existence of what I call value-incommensurabilities. According to Raz (1986), the values of A and B are incommensurate if i) neither A nor B is better or worse than the other and ii) A and B are not of equal (or roughly equal) value.¹²² The value of A and the value of B simply cannot be compared. In the course of his discussion Raz identifies a subclass of incommensurabilities he calls *constitutive incommensurabilities*, which displays the following features. First, if A and B are incommensurabilities of this type, then a person will typically refuse to exchange A for B (when he already possesses A), or B for A (when he already possesses B). Second, the refusal to exchange one option for the other is constitutive of that person's ability to successfully engage in certain pursuits. Finally, and most importantly for our purposes, "it is typical, where [constitutive incommensurabilities] are involved, for agents to regard the very thought that they may be comparable in value as abhorrent" (p. 346). There is, as it were, a mental barrier to exchange, "as if trade-offs involve a heavy price" (ibid). But this attitude is not merely a psychological association to which constitutive incommensurabilities give rise. Rather, it is a qualification for having certain valued relations at all.¹²³

¹²² Raz (1986), Chapter 13. All quotations are from this work. For related material on incommensurability, see Raz, "Incommensurability and Agency" and "Mixing Values" in his (1999).

¹²³ This idea is pursued in a different vein in XXXXX (2004), "Love as a Valuing Relationship" Philosophical Review

This is evident in the example Raz gives of comparing the value of parenthood with the value of making money: “For many, having children does not have a money price because exchanging them for money, whether buying or selling, is inconsistent with a proper appreciation of the value of parenthood” (p. 348). That is, we can attach a value to having children and we can attach a value to one’s earnings, but to ask the question, which has greater value, threatens to undermine one’s ability to recognize what being a parent involves and, for that reason, disqualifies one from entering into a relationship of that sort at all.¹²⁴ The reason, according to Raz, is that

significant social forms, which delineate the basic shape of the projects and relationships which constitute human well-being, depend on a combination of incommensurability with a total refusal even to consider exchanging one incommensurate option for another (p. 348).

Hence, certain relations (e.g. parenthood or companionship) are only available to those individuals whose practical thinking is constrained by an unwillingness to engage in certain value comparisons.¹²⁵

For the purposes of this paper, we can set aside Raz’s central (and quite bold) contention that there exist basic incommensurabilities with respect to value. For even if this claim turns out to be false, it does not threaten the claim that certain relationships depend on a basic reluctance to compare the value of that relationship with various other pursuits. That is, even if those pursuits *can* be compared, it will not take away from the point that the significance of refusing to compare them is

¹²⁴ That one happens to produce offspring may make one a parent in the biological sense, but the assumption here is that parenthood is, at least in part, a normative matter, that is, a matter of understanding which modes of thoughts and action are appropriate to that particular role.

¹²⁵ It is not enough that one act *as if* one is reluctant to make such calculations; it must be a fact about one’s practical thinking that some considerations are simply not “live options.”

integral to one's ability to enter into, and remain in, those relations in the first place. After all, it is one thing to claim, as Raz does, that "*belief* in incommensurability is itself a qualification for having certain relations," and quite another to claim that incommensurability itself or knowledge of such incommensurability is a qualification for having certain relations (p. 351, emphasis added). Also, I will not take up the relationship between social forms and the way in which these confer value onto the various pursuits open to us.¹²⁶ This claim is somewhat orthogonal to our main focus.

Our focus concerns the attitude associated with constitutive incommensurabilities. The general thought is that partly constitutive of parenthood, say, or friendship is the attitude that the value of that relationship cannot be compared to the value of other pursuits; to have such an attitude opens the way to placing a "market price" on that relationship.¹²⁷ Moreover, if one is generally willing to entertain the belief that being a parent or being a friend has its price, that for the right incentive one would be willing to give up one's child or one's friend, then we might reasonably doubt that one recognizes the value of that relationship. Of course circumstances may sometimes force us to reconsider our investments in certain relationships, but this should be kept separate from the disposition I'm drawing attention to, viz., a general willingness to exchange one's relationship for other goods when the opportunity arises.¹²⁸ Such an individual treats his child or friend as a commodity, as something whose value is conditioned on the value of other

¹²⁶ Raz develops this idea further in his (2003).

¹²⁷ This is not to say that one cannot *prefer* one's friend over some other pursuit; as Raz insists, this does not entail that the values are comparable or that one has no reason for valuing the various things one values.

¹²⁸ Thus one can consider the psychological costs of a particularly taxing friendship without indicating a failure to understand the concept of friendship. I deal with this nuance at greater length in the third subsection below.

alternatives. But such treatment reflects a level of disrespect that is inconsistent with the structure of those relationships. It is the mark of friendship, for instance, that one remains loyal to one's friend even when doing so may incur costs that could (conceivably) be dispensed with in favor of having no friend at all. After all, being a friend, like being a parent, is nothing if not a commitment: it is a commitment to recognizing which modes of thought and action are consistent with that role and which modes of thought and action one could not, as a friend or parent, "live with."

It should be noted, however, that one's willingness to make value comparisons where such comparisons are inappropriate does not imply (at least at this stage) a *moral* failure so much as a failure to appreciate important forms of human interaction, where participation in such forms may be part of a flourishing life but are nonetheless morally optional.¹²⁹ It is a tragedy perhaps that one is excluded from entering into, or from remaining in, certain relations, but for all that has been said we need not conclude from this that a moral transgression has occurred. More needs to be said in order to diagnose the sort of failure illustrated in the case of vintage boots. Up to this point, our discussion has taken place in the context of particular relationships like parenthood and friendship. But surely these relationships are optional, at least from the standpoint of what is morally required. My life may be (*may be*) impoverished because I decide not to have children or friends, but there is no reason to think that in so doing I have failed in one of my moral duties.¹³⁰

¹²⁹ The distinction I'm alluding to is the distinction between the obligatory and the permissible, however one wishes to draw this distinction.

¹³⁰ One should count Aristotle among the few if one thinks affability (and the other virtues underpinning friendship) are indispensable virtues.

One way then of connecting Raz's remarks on constitutive incommensurabilities with a set of distinctly moral considerations is by identifying a relationship that goes beyond the sorts of relationships we are (presumably) free to enter into and take ourselves out of. The most obvious candidate would be a moral relationship, a relationship between persons as such. Here it is sufficient that an individual see a reason to treat another in morally decent ways, not out of regard for the particular relation she stands in to this person (e.g. parent, friend, colleague), but out of regard for his personhood. But it is not enough simply to identify another sort of relationship (albeit a universal one) and claim that we have reason to remain loyal to this relation. There is an explanatory condition that must also be met. We must explain what it is about persons as such that makes them the sort of things for which the attitudes associated with constitutive incommensurabilities are appropriate. On this note we turn to Kant.

Dignity and Price

Kant's by-now famous injunction against treating others merely as means is accompanied with the command that we treat others, at the same time, as ends in themselves.¹³¹ According to Kant, persons are ends in themselves by virtue of their humanity, that is, their capacity to set ends. And distinctive of ends in themselves is their *dignity*. Things with dignity, unlike things with mere price, possess "unconditional and incomparable worth." Thus we are obliged to treat others as ends in themselves out of respect for their unconditional and incomparable worth.

¹³¹ Kant (1996): 96 [429].

But it is the distinction between dignity and price that should give us pause.

In his Groundwork for the Metaphysics of Morals Kant claims:

Whatever has a price can be replaced by something else as its equivalent; on the other hand, whatever is above all price, and therefore admits of no equivalent, has a dignity (102 [434-35]).

Thus what possesses price has, in principle, an equivalent with which it can be replaced. But what possesses dignity is priceless: it cannot be replaced by any amount of price. There is, however, an ambiguity in Kant's remarks here that has been the subject of some controversy. Namely, how are we to understand the claim that things with dignity admit of no equivalent? Is this the claim that things with dignity are incomparable in value with things with mere price, leaving open the possibility that, as Thomas E. Hill Jr. puts it, "some things may have more dignity than others, and that the sacrifice of dignity in one sphere might be justified by its enhancement in another" (Hill, 1992, p. 48)? Or, alternatively, is the claim that dignity is *unquantifiable*, such that it makes no sense to say that some things have more dignity than others, thus precluding the possibility of exchanging one amount of dignity for a greater amount of dignity elsewhere?

Hill argues that Kant's remarks imply the latter reading, for had Kant meant that things with dignity admit of no equivalent *in price*, one would have expected him to say so. Similarly, if Kant had meant that the worth of things with dignity is incomparable-to-things-with-mere-price, then Kant misleads us with the claim that things with dignity possess "incomparable" worth. But, for our purposes, we need not decide the matter here. What is plain is that Kant regards the worth of persons as inherently special. How might this bear on our *relationships* to such beings? If I am

engaged in a relationship with a being of special worth, such as a person, then do I show proper respect for that person's worth when my practical thinking is shaped by thoughts of how I might exchange this relationship with other goods, when I am generally disposed to wonder, how much is this relationship worth? This question is of course distinct from the question, how much is this *person* worth. But to regard one's relationship in the way described has direct bearing on the quality of one's regard for the person with whom one is so related.¹³²

Recall that the lesson we took away from Raz was that one's capacity to appreciate fully the value of a particular relationship is conditional on one's "refusal even to consider" exchanging it for another option. For, according to Raz, it is constitutive of some relationships and pursuits that their value is incomparable to the value of other relationships and pursuits. This invites the following thought: If this is the appropriate attitude to take in response to particular relationships like parenthood, in virtue of the incomparable value of that relationship, shouldn't we expect a similar attitude in response to more general relationships among persons, in virtue of the incomparable worth of persons?¹³³ After all, if one threatens the integrity of one's relationship to one's child by considering the "market value" of his life, then does not one likewise threaten the integrity of one's relationship to another person by

¹³² Darwall (2004) has recently offered a complementary account of Kant that emphasizes the "second-personal" reasons for respect, where a second-personal reason "is one whose validity depends on presupposed authority (hence accountability) relations between persons and, therefore, on the possibility of the reason's being *addressed* person-to-person" (p. 46).

¹³³ We must keep separate the incomparable worth, if any, of relationships with the incomparable worth of persons. There may be exceptions to the idea that the value of a valuable relationship is inherited from the value of the individual to whom one is related, but I am claiming this is the case in the moral realm.

considering the “market value” of her life? We do not treat our children like commodities precisely because we could not continue to be parents if we did.¹³⁴

Again, this line of thought is not held hostage to the possibility that, as a matter of fact, the worth of persons, like the worth of one’s relationship to one’s child, *can* be compared to the worth of other pursuits. For what we saw in the previous subsection is that it is the *regard* one has for one’s child or one’s friend that matters from the perspective of that relationship. Loyalty or fidelity to that relationship is decoupled from the possibility that “one could do better” if one chose another course. Whether or not one could is beside the point. The point is that many of the relationships we care about—including what we might take to be the most basic relationship of all, viz., the relationship among persons in a moral community—depend on a mode of practical thinking that respects the “non-exchangeable” or “non-fungible” worth of each individual. A friend who is generally open to the possibility of “trading-up” is no friend at all.¹³⁵ Likewise, an agent who is open to the possibility of exchanging someone’s worth for the worth of someone else exhibits a fundamental disrespect for the moral relationship one bears to that person.

In sum, Kant’s thesis emphasizes the intrinsic worth of persons. But Kant-plus-Raz emphasizes the *judgment that* persons are of intrinsic worth, since it is this judgment (or, more generally, this attitude) that is part of the grounds for moral

¹³⁴ Sam Kerstein has reminded me that in many places of the world, places where living conditions are perilous, parents may well have to treat their children like commodities, but where this does *not* imply that these individuals are incapable of appreciating the value of parenthood. For example, The New York Times (June 23, 2003 (A3)) recently ran a story about mothers and fathers in India having to sell a child in order to feed their remaining children. Examples such as these only reinforce the point made above that absolute prohibitions against comparisons are untenable. See below for further discussion.

¹³⁵ Construing morality in terms of an ideal of human relationships (like friendship) has found favor with a number of contemporary non-consequentialists; in section II below, I offer one such construal as a way of showing how my account of unique dependency is precisely the sort of view one should hold to the extent that one is drawn to non-consequentialism.

relationships in the first place. Indeed, one might argue, though I won't here, that this attitude is part of our most basic moral attitude.

Unique Dependency

It is but a short step from here to the claim that unique dependency possesses distinct moral significance. Recall **Vintage Boots**. Peter is confronted with a child whose life is directly threatened by harm. But Peter is also aware that he is the only means of removing that harm; without his help, the child will soon drown. For this reason, we say that the child is uniquely dependent on Peter. But Peter chooses to ignore the child's pleas on the grounds that by doing so he can prevent more harm from occurring abroad than he can prevent here and now.

When first confronted with this case, most of us find Peter's reasoning deeply troubling; not only are we troubled by his decision to refuse to rescue the child, but we're troubled by the reasons he cites in favor of that decision. It is time now to give content to this worry.

As noted at the beginning of this section, what stands out most about **Vintage Boots** is Peter's willingness to contemplate the relative values of rescuing the child versus selling the boots and saving distant others. We can now explain why this feature should be of moral interest to us. Peter regards the value of this child's life as commensurate with the value of other lives abroad, since he explicitly compares the disvalue this child's death with the disvalue of several deaths abroad. But it is precisely this willingness to engage in such comparisons that reflects a failure to appreciate the value of his relationship to the child, a relationship marked by unique

dependence. There is more than a surface similarity between a parent's abhorrence at the thought of exchanging her child for money and the abhorrence most of us experience at the thought of refusing to rescue the child for the sake of starving individuals abroad. This does not mean that we see no reason to respond to the suffering of distant individuals; surely we do. Analogously, a parent's abhorrence at the thought of exchanging her child for money does not entail that she see no reason to increase her wealth.

The point is that parents display a general and abiding reluctance even to consider exchanging one for the other because this is part of the normative standard for being a parent. Correlatively, moral agents display *inter alia* a general and abiding reluctance even to consider exchanging the life of someone uniquely dependent on them for other lives elsewhere because this is part of the normative standard for being a moral agent. We do not—indeed, should not—regard the lives of other people as transferable goods.¹³⁶ But why not think that the normative standard for being a moral agent involves saving the greatest number of lives? Am I not simply begging the question against the consequentialist?

I would indeed be guilty of such a move if I were merely assuming that part of the moral standard for being a moral agent involves a general and abiding reluctance to consider such comparisons. But I have tried to mount an argument for this claim, based in part on shared (or what appear to be shared) intuitions in response to **Vintage Boots**. If the normative standard for being a moral agent indeed involved

¹³⁶ After all, few, if any, of us regard our own lives as such. That is, we do not regard our own worth as contingent on how it might be exchanged for an increase in worth elsewhere, *even if* there is some metric according to which our worth can be measured against the worth of others. For we have, at base, a reason to be “true to ourselves,” and this reason is not reducible to reasons of general welfare promotion. See Raz (1986), Chapters 13-14; Hampton (1993); and Wolf (1982).

saving the greatest number of lives, then we should be surprised by our considerable reluctance to rescue the drowning child. Without obvious reasons for dismissing this intuition, I have argued that what best *explains* this reluctance (and the reluctance we experience in similar cases¹³⁷) has to do with a more general attitude that is partly constitutive of being a moral agent. A defender of this rival standard needs a way of accounting for or, at the very least, explaining away this reluctance.

One standard consequentialist tactic is to move to an indirect or two-level account, according to which there are compelling consequentialistic reasons why one ought to adopt or preserve certain apparently non-consequentialistic attitudes and dispositions.¹³⁸ In the present case, the claim would be that adopting or preserving the reluctance I've identified would, under standard circumstances, yield better consequences than either adopting or preserving alternative attitudes, or simply disabusing oneself of the reluctance altogether. But this move is unsatisfactory for several reasons. First, as an empirical matter, it seems far more likely that we would approximate the normative standard of saving the greatest number of lives if we succeeded in suppressing, if not disabusing ourselves of, such a reluctance. Second, by pressing for the inclusion of this reluctance, the indirect consequentialist risks conceding the very point she wishes to reject, viz., that this reluctance is not morally dispensable. For if the repugnance we feel towards an individual that lacked such a reluctance is supposed to be part of an explanation of why we should resist the suppression of such attitudes, then surely this indicates that, as an empirical fact, we

¹³⁷ See **Detention Center I and II**, Chapter 3.

¹³⁸ We considered several of these views in Chapter 4.

take this reluctance to be a basic part of our moral personality, *independent of its role in utility calculations*.¹³⁹

Of course one might reasonably wonder how rigorous this inhibitory attitude ought to be. Is it, for example, constitutive of one's capacity to enter into moral relationships that one refuse even to consider the comparable worth of saving one child from drowning versus saving a *thousand* children from drowning? Surely one does have a reason to consider such an exchange. But no one would suggest that one's moral credentials are open to question as a result. Indeed, the reverse seems to be the case: reluctance to consider such an exchange reflects a rather dangerous "moral provincialism."

But this objection trades on a mistaken assumption, viz., that one's moral relations to others *absolutely prohibits* comparative assessments. The thesis, rather, is that one's moral relations to others should be marked by a "general and abiding" reluctance to make such comparative assessments. By way of analogy, consider how a parent of four, one of whom is severely handicapped, might reason. She might consider devoting more time and resources to the handicapped child, knowing full well that this will deprive her other children of what they reasonably deserve. Now according to my thesis we need not conclude that her credentials as a parent are diminished as a result of her reasoning, unless of course it comes to light that all (or even a preponderance) of her decisions about how to treat her children are marked by an insensitivity to the special worth of each of her children. We might suspect such insensitivity, for example, if her decision bore no special marks of emotional turmoil,

¹³⁹ But won't utility calculations nevertheless always "outweigh" considerations of unique dependency? No, for there are other ways individuals will be uniquely dependent on us, ways that do not directly concern one's biological needs.

at least in the early stages of her decision. In any event, if we are prepared to accept these remarks, then it is possible to see how our moral relations to others might, under certain circumstances, allow for comparisons. Like much else in moral philosophy, however, there is unlikely to be a principled threshold according to which our attitudes suddenly take account of “the numbers.”

Allegiance to the idea of unique dependence is, first and foremost, an allegiance to a mode of practical thinking. It is a mistake, however, to infer from this that difficult questions cannot penetrate this reluctance, just as difficult questions cannot penetrate a parent’s reluctance to compare the worth of her children. To pinpoint the difference between Peter and our response to the dire cases, it is enough to identify the differences in our general starting points: Peter, wedded to his humanitarian cause, begins with an attitude congenial to comparison since that cause amounts to maximizing the prevention of suffering. It is only when the fact of the child’s unique dependence is thrust upon him that he is prepared to consider (though only consider) his reasons for ignoring his humanitarian ends and saving the child’s life.¹⁴⁰ By contrast, most of us begin with the attitude that someone’s unique dependence on us is a compelling reason to rescue him, independent of its bearing on the overall state of the world. It is only when the prospect of catastrophe is thrust upon us that we are prepared to consider our reasons for ignoring the child and averting the catastrophe.

¹⁴⁰ I grant that there may be other descriptions of Peter’s reasoning, one’s which involve different attitudes toward the relevant parties. But insofar as an individual reasons along the lines I have described, insofar as he “leads” with his humanitarian concern come-what-may, I stand by the assessment in the text.

II. An Ideal of Human Relationships

I have argued that moral agents display *inter alia* a general and abiding reluctance even to consider exchanging the life of someone uniquely dependent on them for other lives elsewhere because this is part of the normative standard for being a moral agent. But how is this standard related to the rest of morality? It might be thought that I am promoting a sort of perfectionism, what with this emphasis on attitudes and dispositions. But this is not the picture I wish to convey. The reluctance I've tried to spell out is connected to a more fundamental conception of morality, one that emphasizes an *ideal of human relationships*.

Part of the task for moral philosophers, on my view, is to identify the features of this ideal and how these features place obligations and constraints on our conduct.¹⁴¹ Much of what I have tried to do in this work is to illuminate one part of this connection. In Chapters 2 and 3, I argued that one of the central difficulties for utilitarian accounts lies in their inability to explain the wrongness of refusing to rescue someone in paradigm cases of rescue. When we consider the cases carefully, it is apparent that what the utilitarian overlooks is *the direct way in which I wrong the victim* when I refuse to act. I do something to him that is independent of the “overall state of affairs” (assuming we know what this means). And what this is can best be characterized in terms of my relation to the victim: the consideration that counts in favor of my rescuing him (viz., the fact that he is uniquely dependent on me) also grounds his claim on me, since he is especially vulnerable to being wronged by me. He has, as R. Jay Wallace puts it, a “privileged basis for moral complaint, resentment,

¹⁴¹ We should not expect, however, that the features of this ideal will coalesce into a neat picture where *inter alia* values are hierarchically arranged. That we enter into many varieties of relationships with others is an important fact in evaluating the varieties of reasons we have for responding to need.

and so on, precisely insofar as I have acted with indifference to the value of relating to [him] on a basis of mutual recognition and regard.”¹⁴² Thus the reluctance that is part of the normative standard of moral agency is grounded in a more basic concern for the value of our relations with others. And, as I’ll argue below, such a concern manifests itself in the way in which we strive to conduct ourselves in ways others could accept—or, more specifically, in ways that no one could reasonably reject.

That morality is modeled on an ideal of human relationships is what we should expect if value has the character I described in the previous chapter. I argued there that our reason to respond to the suffering of others does not flow from the impersonal disvalue of suffering, which is itself a consequence of the idea that suffering is bad in itself. I tried to show that this idea does not in the end make sense. A more plausible idea is that suffering is bad because it is bad *for* someone, because someone typically has reason for wanting to avoid suffering. So it is that we have reason to respond not to a person’s ends as such (after all, she may have reason to want to endure suffering), but to the person herself. But we need to move beyond these rather suggestive ideas to something a bit more specific: how does the relational conception of morality manifest itself in our practical affairs? That is, how is this supposed to constrain and shape our practical thinking? The most promising answer, in my view, is offered by contemporary contractualist views. What follows is not an attempt to provide a thoroughgoing defense of this view; instead, I’ll argue that the contractualist framework offers a plausible *realization* of the idea that morality is essentially a matter of how we relate to others.

¹⁴² Wallace (unpublished), p. 35.

For the contractualist, the reason-giving force of moral considerations stems from an agent's recognition of a particular value; on Scanlon's (1998) contractualist account, that value consists of living with others on terms that they could not reasonably reject. We are, on this picture, "answerable to others," not simply in the punitive sense that we can be sanctioned for wrongdoing, but rather in the deliberative sense that we seek to shape our conduct in ways that others could accept. Quite obviously, a consequentialist conception need not commit itself to such an ideal. But the value of living with others on terms they could not reasonably reject is actually multi-dimensional. To say that Scanlon has identified *one* particular value is actually misleading, since the value in question appears to have its roots in several values. There is the general value of being in relationships to others, the value of justifying ourselves to others, and finally the value of shaping our practical reasoning in accord with others'. The reason-giving force of moral considerations, then, flows from a recognition of these mutually-reinforcing values.

When, therefore, we reflect on the general and abiding reluctance to consider certain value comparisons, that which I described as part of the normative standard for moral agency, we can begin to see how this reluctance fits into the Scanlonian model. Someone who is generally open to evaluating the relative values of different relationships, and who is generally prepared to exchange his relationships for other goods, does not make himself answerable to others; instead, he is answerable either to certain personal ideals or, as in the case of the utilitarian, to a maximizing calculus. In either case, such a person *transcends* the interpersonal ideal of justifiability in a way that we should find troubling. For what the Scanlonian model keeps firmly in

place are the *individuals* who must live with our actions, and so it is crucial that we direct our moral consultation at the individuals involved rather than at distributions of goodness across possible states of affairs. Insisting on a “centralized” notion of good states of affairs and imposing this on others is, at best, a stilted way of relating to others. At worst, it undermines the relation of mutual recognition.

If we are prepared to accept this idea of mutual recognition and the importance of our relationships to others, it will come as little surprise that the practical realization of these ideas consists in acting on principles no one similarly motivated could reasonably reject. The ideal of mutual recognition and respect leads naturally to a distinct ideal, viz., the ideal of justifiability. The thought is that recognizing the value of other persons and our relationship to them places a critical constraint on our conduct in this sense: we must refrain from acting on principles specifying general rules of conduct that others could reasonably reject. What might the grounds for rejection be? They might be any number of things, and there is no need to specify these in advance. The point is that the reasons for rejecting principles are not themselves constitutive of the *wrongness* of one’s action, even though such reasons figure centrally in the rejection of a given principle of conduct. On the contractalist’s account, it is the fact that others could reasonably reject the principle on which one acts that makes an act wrong.¹⁴³ And this underlines the relational aspect of morality at the heart of contractualism.

¹⁴³ Pursuing this a bit further, we should see that even a consequentialist can accept the biconditional that an action is wrong if and only if it violates a principle no one could reasonably reject. What a consequentialist should reject is the stronger claim that an action is wrong *because* it violates a principle no one could reasonably reject.

Putting these points together, we can now explicate how it is that I directly wrong the individual uniquely dependent on me, something the consequentialist was unable to explicate. Because he is uniquely dependent on me, he has a conclusive reason to want my assistance, since I am the only individual in a position to do so. The relevant moral question, for the consequentialist, would be: how would the consequences of assisting him compare to the consequences of not assisting him? For the contractualist the question is: in the face of his reason for wanting my assistance, do I have reasonable grounds for rejecting a principle requiring me to assist him? Or, more simply, could I justify not rescuing him? It may be observed that rescuing him may be inconvenient or messy or emotionally taxing, but none of these considerations constitute reasonable grounds for rejecting the relevant principle—which, it should be emphasized, is grounded on *his reasons* for wanting to be rescued. In short, I cannot justify to him my refusal to act. Thus, if I refuse to rescue him, I violate a rule of conduct I myself could not reject.

What morality requires of me then is that my actions be justifiable—not in light of some abstract standard, but to whomever my conduct directly impacts. This is why contractualism, unlike its main normative rival consequentialism, is said to illuminate the relational aspect of morality. But this last point has important implications for principles of humanitarianism, for the content of the drowning individual's reasons for wanting assistance is fundamentally distinct from the content of the reasons individuals abroad have for wanting assistance. This is a question I take up in the next chapter. In the remainder of the chapter, I want to address

potential problems with my account of unique dependence and the demands such dependence places on potential rescuers.

III. Dependency and Responsibility

How seriously should we take this notion of unique dependency? For surely we can conceive of scenarios in which you have an obligation to rescue X *despite the fact* that X is not uniquely dependent on you. Picture, for example, a beach crowded with able-bodied swimmers. Are we to assume that when you and the other swimmers become aware of a child drowning in the shallow surf, you have no duty to save him? After all, he is not uniquely dependent on you. For that matter, since he is not uniquely dependent on any one of the swimmers, it would seem to follow that *no one* has a duty to rescue him. This is highly implausible. Perhaps then unique dependency is a distinction without a moral difference.

In order to approach this concern, I want to shift the terms of the discussion just slightly and consider the matter from the point of view of responsibility. I will argue that in the case above (call it **Crowded Beach**) you are responsible, not for the outcome (*viz.*, the drowning death of the child, assuming he is left to drown), but for your failure to make an effort to rescue him. For what caused the child's death was the *collective* failure of the group, and that is something for which you cannot be responsible. By contrast, Peter *is* responsible for the outcome in **Vintage Boots** since he caused that outcome.¹⁴⁴ Appreciating the difference requires some explaining.

For each of our two cases (**Crowded Beach** and **Vintage Boots**), I want to ask: first, what caused the outcome? And second, who, if anyone, is responsible for

¹⁴⁴ For the sake of the argument, I am assuming that one can cause outcomes by omission.

the outcome?¹⁴⁵ In **Vintage Boots**, the answers are fairly straightforward. The child's death is caused by Peter's refusal to rescue the child—again, assuming that omissions can be causes. If $F(P)$ represents Peter's failure to take the necessary steps to prevent the child's drowning, then $F(P)$ is a sufficient cause of the child's death.¹⁴⁶ It is thus natural (and I believe correct) to hold Peter responsible *for the child's death*. Since Peter's omission was a sufficient cause of the child's death, we rightly hold Peter responsible for the death.

But matters are not the same in **Crowded Beach**. Your failure to respond is not the cause of the child's death, even if it is *a* cause. Let me explain. If $F(Y)$ represents your failure to take the necessary steps to prevent the child's drowning, pretty clearly $F(Y)$ is not a sufficient cause for the child's dying, since the child is rescued just in case at least one other swimmer takes action. Perhaps one swimmer acts before anyone else has a chance to. So if we suppose there are five able-bodied swimmers, then the child is rescued if, but not only if, $F(Y) \wedge (A \vee B \vee C \vee D)$ obtains. Any individual failure is not sufficient for death; indeed, the failure of four-fifths of the group is not sufficient. What *is* sufficient for the child's death is a case in which $F(Y \vee A \vee B \vee C \vee D)$ obtains, which is equivalent to $(F(Y) \wedge F(A) \wedge F(B) \wedge F(C) \wedge F(D))$. Hence, what causes the outcome in **Crowded Beach**, where the outcome is the child's death, is $F(Y \vee A \vee B \vee C \vee D)$. Strictly speaking, then, it is false to say that you cause the child's death; collective failure causes his death. And

¹⁴⁵ The following discussion has been greatly influenced by Sartorio (2004). Many thanks to Mark Schroeder for bringing it to my attention.

¹⁴⁶ For the purposes of discussion, I am ignoring cases of moral luck. As an example of moral bad luck, the child dies (he's attacked by a shark) even though you swim out to rescue him; as an example of moral good luck, the child is rescued (a wave tosses him to the shore) even though no one attempts a rescue.

this is perfectly compatible with saying that if, but not only if, *you* had taken the necessary steps, his death would have been prevented. At any rate, **Crowded Beach** is a standard case of *collective causation*. So what then should we say about responsibility: who, if anyone, is responsible for the outcome?

The obvious, and I believe correct, answer is: *you all* are responsible. There is collective responsibility for the child's death. But then how should we understand the distribution of responsibility among members of the group? In what sense, if any, are *you* responsible? I submit that you are responsible for the obtaining of $F(Y)$, but you are not responsible for the obtaining of $F(Y \vee A \vee B \vee C \vee D)$, since this is not something you could be responsible for.¹⁴⁷ You are of course responsible for making an effort to rescue the child, and failure to do is reason for serious moral condemnation. But since $F(Y)$ does not entail $F(Y \vee A \vee B \vee C \vee D)$ and since $F(Y \vee A \vee B \vee C \vee D)$ is not only sufficient for the child's death but, barring moral good luck, necessary, you—and you alone—cannot be responsible for the child's death. After all, you cannot be responsible for something you could not be responsible for, viz., $F(Y \vee A \vee B \vee C \vee D)$. It would be puzzling if this were not the case.

The contrast with **Vintage Boots** then should be clear: there, $F(P)$ is a sufficient cause of the child's death and since Peter can be responsible for $F(P)$, Peter is directly and solely responsible for the child's death. Not so in **Crowded Beach**, as we've just seen. But one may nevertheless maintain that despite this difference, you have just as much reason to rescue the child in **Crowded Beach** as Peter does in

¹⁴⁷ Strictly speaking of course, this is false. You could threaten the other swimmers with harm if they attempt a rescue; you could assure them that *you* will attempt a rescue without following through; and so on.

Vintage Boots. To test this claim, we'll need a third case, one which stands between **Crowded Beach** and **Vintage Boots**.

Call the following case **Two Pools**. You happen upon two pools. In the West Pool, some fifteen meters to your left, is a small child evidently in distress. In the East Pool, some fifteen meters to your right, is another child similarly in distress. However, around the East Pool stand four apparently healthy young men and women, gathering at the pool's edge.

Quite clearly, you should attempt to rescue the child in the West Pool, where your help is the only help available. Indeed, you reasonably assume that at least one of those four observers will attempt to rescue the child in East Pool. You take it for granted, that is, that $F(A \vee B \vee C \vee D)$ will not obtain—that $F(Y) \wedge (A \vee B \vee C \vee D)$ will. Still, you have some reason to rescue the child in the East Pool, since you have reason to avoid $F(Y \vee A \vee B \vee C \vee D)$, which is a sufficient cause of death. But does it seem at all plausible to assert that, on these grounds, one has just as much reason to go to the East Pool as to the West Pool? True, one has just as much reason to keep $F(Y \vee A \vee B \vee C \vee D)$ from obtaining at the East Pool as $F(Y)$ at the West Pool; after all, both deaths are equally tragic. But this obscures the point: the point is that you have every reason to assume that at least one member of the group at the East Pool will act, and if they do not act and $F(Y \vee A \vee B \vee C \vee D)$ obtains, you can justifiably claim that while you are indeed responsible for $F(Y)$ at the East Pool, you are not responsible for $F(A \vee B \vee C \vee D)$.¹⁴⁸ Moreover, whereas $F(Y)$ is not a sufficient

¹⁴⁸ Suppose, however, that you know that $F(A \vee B \vee C \vee D)$ will obtain, because, for example, the onlookers are evidently disabled. Thus stipulated, the case is straightforward: flip a coin. Matters are not so straightforward when you have good reason to believe that $F(A \vee B \vee C \vee D)$ will obtain,

cause of the child's death at the East Pool, F(Y) *is* a sufficient cause of death at the West Pool. It seems plain from **Two Pools** then that there is a reason to respond to an individual uniquely dependent on you that is weightier than the reason to respond to an individual not so dependent.¹⁴⁹ And claiming this does not in anyway entail that you have little or no reason to respond to individuals who are not so dependent on you, as in **Crowded Beach**.

In general then, I see no reason to accept the claim that the *only* relevant consideration in virtue of which you have a reason to aid others is the impersonal badness of suffering. It is false, in other words, that the wrongness of your refusal to act beneficently can be accounted for entirely by reference to the suffering you refused to prevent. Unique dependency makes important differences in both the stringency of our reasons to aid others and the sorts of wrongs we perpetrate when we fail to act.¹⁵⁰

Where my argument departs from previous attempts to separate distant suffering from local emergencies is in suggesting that these reasons do not even *begin* in the same place. It has long been assumed that our obligation to support humanitarian relief efforts is simply an extension of our obligation to intervene in cases exemplified by the drowning child. It is only when we get to the details of the two cases that distinctions are drawn. By contrast, I have argued that unique dependency changes the set of considerations that bear on moral deliberation. If

because, for example, the onlookers are laughing and yelling racial epithets at the child. The difficult questions concern responsibility and the ways in which others' omissions can affect your responsibility. I will take up these questions in the next chapter.

¹⁴⁹ Put differently, there is a difference, then, in the stringency of one's reason to respond when one believes, on the one hand, that that child needs someone's help and, on the other, that that child needs someone's help *and I'm the only one who can help him*.

¹⁵⁰ That 'determinacy' can likewise make a moral difference is conclusion of my "Benefits without Beneficiaries," under review.

nothing else, it reveals the mistake in Singer's claim that "we have an obligation to help those in absolute poverty that is *no less strong* than our obligation to rescue a drowning child from a pond" (Singer, 1993, p. 230, my emphasis). Regardless then of the fact that we only rarely, if ever, stand in the relation of unique dependency to others, this relation occupies a different place in our moral thinking than does our relation to distant sufferers, and so we distort our understanding of our relation to this latter group by modeling it on our relation to the former group. This result, I contend, is worth elucidating.

There is, however, a further story to be told about our responsibility toward those in the greatest need. In this chapter, I have argued that when the relation of unique dependency obtains, the donor has a perfect duty to respond, a duty that is connected to the special relationship between donor and recipient. But no such relation exists in the case of humanitarian relief: there is no (determinate) individual that relies on *you and only you* for relief from suffering. So how does this bear on our moral obligations to contribute to humanitarian relief efforts? In the next chapter, I will argue that obligations to eradicate world suffering fall principally on collectives, like the governments of developed nations, for while a potential recipient of humanitarian aid *H* cannot reasonably reject a principle allowing *you* to pursue a wide range of projects that contribute nothing to humanitarian relief, *H* can reasonably reject a principle allowing an entity like the US government to take an active role in eradicating world suffering.

CHAPTER SEVEN

Humanitarianism

IF YOU ARE MORALLY REQUIRED to respond to the needs of an individual uniquely dependent on you (because she is uniquely dependent on you), what, if anything, are you required to do for those individuals who are not uniquely dependent on you? This is a slightly different question from the question we confronted towards the end of the previous chapter: there, the question was whether reasons to respond to individuals uniquely dependent on you are *more stringent* than reasons to respond to individuals not so dependent. I argued that they are. But answering that question leaves open what requirements, if any, you might have toward individuals who do not uniquely depend on you for rescue. In the case of **Crowded Beach**, I argued that you are required to respond to the needs of the drowning child, even though he is not uniquely dependent on you. How then can we explain this? And how will this explanation shed light on our moral relationship to those living in extreme poverty abroad?

The present chapter provides answers to these questions. First, I will argue that unique dependency can be extended quite naturally to cover cases in which there are multiple donors, as in **Crowded Beach**. While the drowning child is not uniquely dependent on *you*, he is uniquely dependent on the five of you, since the group's failure to mount a response (either implicitly or explicitly) practically guarantees the child's imminent death. Second, this account can also be extended in such a way as to illuminate the puzzle of humanitarian relief: it can be argued that while individuals

living in extreme poverty are not uniquely dependent on you or me, they are uniquely dependent on groups of which you and I are members, for example, a wealthy nation like the US. Finally, I show how this account follows from the account of moral justification implicit in the contractualist framework discussed in the previous chapter.

I. One-to-Group Dependence

How do your moral reasons to provide aid change when there is more than one donor capable of providing aid? Let us assume that there is a determinate individual, e.g. a drowning child, who stands to suffer a serious loss unless he is rescued. (We will consider cases in which potential beneficiaries are indeterminate when we move to a discussion of humanitarian relief in Sections IV and V.) In cases of what I'll call *one-to-group dependence* the group has an obligation to respond to the child, an obligation that is different in character than your obligation to respond to an individual uniquely dependent on you. This is because the child is uniquely dependent on the group, not on you. After all, it's obvious to all that *someone* ought to rescue the child—someone, that is, who is in a position to rescue the child. But since, by hypothesis, there are a number of individuals in a position to rescue the child, we do not want to say that any *and all* individuals in a position to rescue the child ought to, for there are practical considerations that count against *everyone's* attempting a rescue. Quite obviously, if everyone dives into the surf after the child, this could undermine the rescue effort in a way that a coordinated effort would not. We do not want to multiply the number of emergencies by acting recklessly.

What we should say, therefore, is that the group of individuals in a position to rescue the child ought to rescue the child through coordinated action. By ‘coordinated action,’ I mean simply that the group should discharge its obligation in whatever way will bring about the desired end. In some instances, coordination will amount to little more than that which occurs in a baseball game when one player “calls off” his teammates from a fly-ball he intends to catch. In more complicated instances, coordination will involve delegating out different responsibilities (calling 911, performing CPR, etc.). In the most complicated of instances, viz., humanitarian relief, coordination will involve multiple specialists (ecologists, epidemiologists, economists, health specialists) working through multiple agencies at multiple levels. But the point is that the group represents an entity in its own right, insofar as rescue is concerned: while there may be latitude in the way that responsibility is distributed among members of the group in one-to-group dependence since there are multiple ways in which the group can discharge its duty, *there is no latitude at the level of the group*.¹⁵¹ The group has a moral obligation to respond to the individual in need. And the reason is that that individual is uniquely dependent on the group.¹⁵²

One of the implications of this approach is that there may be reasons not to respond to the child that are absent in cases of one-to-one dependence (cases like **Vintage Boots**), the most obvious case being one in which someone else is already providing assistance. But there are also cases in which one has committed oneself to other causes (like the extreme case depicted in **Two Pools**). At any rate, what is clear

¹⁵¹ In this regard, if we treat the group as an individual, we have in effect a one-to-one dependence.

¹⁵² But what if coordination among members of the group is impractical, if not impossible—as was the situation in the case of Kitty Genovese? Such cases will be dealt with in Section II below.

is that one's moral relationship to the individual in need is mediated by the group and the group's control over the child's fate.

This approach, I contend, is attractive for several reasons. First, by coordinating our responsibilities with the group, we can prevent unintended emergencies that arise when individuals act autonomously. It is sometimes the case that multiple donors, acting autonomously, decrease the chances of rescue—if, for example, the victim's condition is especially precarious. More frequent are cases in which the autonomous actions of multiple donors actually increase the number of emergencies, as when the donors themselves become imperiled.

This points to a second virtue of group coordination: appropriate division of labor. As the nature of a given rescue operation becomes more and more complicated (with humanitarian relief being the most complicated), groups have distinctive advantages over individuals: for example, groups can draw on the special skills and capacities of its members in directing aid. But the most noteworthy advantage of groups over individuals is the creation and deployment of *institutions* that fill the gaps where individuals cannot.¹⁵³ Examples range from fire and rescue to the World Bank. The existence of such institutions fundamentally changes our relationship to individuals in need by providing a critical mediating force. And the very presence of this mediating force is morally significant, for when emergency situations arise, often the most responsible thing a person can do is notify *others*, for example, emergency responders. Intervening oneself is sometimes not simply impractical, but irresponsible. For example, it would be irresponsible of me to confront a house fire with buckets of water *instead of* alerting the fire department. If you are unfamiliar

¹⁵³ See O'Neill (1986), Chapter 3.

with the basics of CPR, it would be irresponsible of you to attempt such a procedure instead of contacting individuals who can.

It would be difficult to overstate this point, and yet it is frequently overlooked in discussions of rescue: we have a moral responsibility to be aware of whether and how one's individual assistance ought to be coordinated with that of the group. In practice, this will mean asking, first and foremost, what does the imperiled individual need from us (whether from someone or from all of us together)? And, second, what does the group need from me? Sometimes, the needs of the individual are such that the group requires that you stay out of the way, as when an ambulance team arrives at the scene of an accident. Other times, the needs of the individual requires that *you* swim out to the drowning child, as when you are the only beachgoer who has not been drinking.

It might be wondered, however, whether extending unique dependency to groups is compatible with the moral foundations for unique dependency spelled out in the previous chapter. Recall that what makes unique dependency a compelling moral reason for action has to do with the dignity of persons, a property which places a barrier on comparing the worth of one individual with the worth of other persons or pursuits. That an individual is uniquely dependent on me is a reason for me to exclude considerations of exchange. But do these sorts of considerations extend so easily to cases of one-to-group dependence? There are no obvious obstacles at least to such an extension. That an individual is dependent on a group, as opposed to an individual, does not seem to make any obvious difference to the moral worth of the imperiled individual; after all, if an individual possesses dignity simply in virtue of

being a person, surely this is unaffected by the number of individuals on whom that individual depends. And it seems fair to say that the group, considered as an entity, ought to be reluctant to compare the value of this individual with the value of other persons or pursuits. It is this last point, however, that presents a puzzle.

When you step back and consider the variety of groups of which you are a part, it seems reasonable to suppose that there are, at any given time, a number of individuals who are uniquely dependent on those groups. So each of these groups is prohibited from comparing the worth of the individuals uniquely dependent on them. And as a member of several of these groups, you must not only recognize the dignity of these separate individuals, but act in ways that correspond to that dignity. But this path is headed for a dilemma. For in acting so as to respect the dignity of one individual, you must exclude from consideration the needs of another individual, but surely your excluding the needs of another individual does not respect the dignity of that individual, who as it stands is uniquely dependent on another group of which you are also a member. Thus it seems that extending unique dependency to groups will inevitably lead to practical dilemmas.

This objection points to a number of ambiguities in the way one-to-group dependence was laid out. Once those ambiguities have been addressed, we will see that this objection is little more than a call for clarification. When an individual *S* is uniquely dependent on a group, the group assumes an obligation to treat *S* as a distinct source of value, distinct, that is, from the other sources of value to which it might respond. It would be wrong, in other words, to treat *S*'s suffering as interchangeable with other instances of suffering—so long as the group represents the

only entity that can prevent, at little cost or risk to itself, serious harm from coming to *S*. But while the group is constrained in the way it responds to *S*'s plight, any one member of the group is not, for there are (let us assume) multiple ways in which the group can discharge its duty. If so, then even if we assume that the group has the responsibility to provide assistance to *S*, a responsibility that cannot be exchanged for other acts of beneficence, it is an open question how the members of the group ought to carry out its responsibility. It may well be necessary to introduce other considerations into deliberations. And there is no obvious reason why these considerations must be bound by the constraints mentioned above, since it is agreed that the group does not seek to compare the disvalue of *S*'s suffering with other sources of disvalue.

Suppose that I am among the five people, say, who observe the small child drowning in the shallow surf (as described in **Crowded Beach**). Since the child is uniquely dependent on the five of us, since our joint failure to render assistance practically guarantees his death, we *as a group* have a perfect duty to render rescue. But since it is not necessary—nor responsible—for all of us to enter the water, a decision needs to be made, assuming (unrealistically) that no one has already entered the water. Suppose that I am looking after my two-year-old son. Would it be unreasonable of me to consider my paternal duties in deciding to what extent, if any, I should participate in rendering rescue? Surely I do not risk the moral standing of my relationship with the drowning child if I consider the welfare of my own son. Perhaps what I should do, if it becomes increasingly clear that none of the other members

intends to render rescue, is to instruct one of the other members to look after my son as I enter the water.

The point to bear in mind, however, is that even if the group is bound by the moral obligation to render rescue in virtue of its relation to the drowning child, there is a question that arises in the case of one-to-group dependence that does not arise in cases of one-to-one dependence, a question about one's supporting role in a group's coordinated action. And asking this question does not amount to asking, as in the case of one-to-one dependence, how the disvalue of this child's death compares to the disvalue of other projects or relations to which one might otherwise respond.

So far, I have been assuming that group membership is unambiguous: in **Crowded Beach**, the group of individuals aware of the drowning child is stipulated at the outset. But how do your moral obligations change when it is unclear *to you* whether the imperiled individual is uniquely dependent on you or uniquely dependent on a group, of which you are a member? The case of Kitty Genovese is representative: Ms. Genovese was assaulted in the courtyard of an apartment building where the number of individuals aware of her distress could vary wildly. Suppose that you hear Ms. Genovese's cries for help. Although it is late at night, it occurs to you that others may hear her cries as well, but you cannot be sure. It seems plain that you have a reason to take the minimal steps necessary to respond to the crisis. Pretty clearly, coordinated action is not an option here; you are precluded from asking what your supporting role to the group ought to be since, as far as you know, you are the only member of the group. At the same time, we can distinguish between what you

have reason to do (viz., rendering assistance) and what you are subsequently responsible for. This builds on remarks made in the previous chapter.

As a normative matter, you have a compelling moral reason to take the minimal steps necessary to assist Ms. Genovese whether or not you are involved in a one-to-one or one-to-group dependence: refusal to do so would reflect a failure to recognize the dignity of her as a person. Besides, not only are there no costs to you of assisting her, but there are no coordination problems that could conceivably oppose acting—we are talking about calling 911, after all. At the same time, in the event that Ms. Genovese is uniquely dependent on you and you refuse to assist her, then your responsibility goes beyond your responsibility for refusing to assist her, which is what you would be responsible for in the event that you refused to respond in a case of one-to-group responsibility. In the former case, you are solely responsible for allowing the harm to come to her.¹⁵⁴ For a sufficient condition for Ms. Genovese's death in a case of one-to-one dependence is your refusal to assist her, whereas such a refusal is not sufficient in the case of one-to-group dependence, as we saw in the previous chapter. Joint refusal is necessary in order to guarantee Ms. Genovese's death, but this is something for which you cannot reasonably be held accountable.¹⁵⁵ Thus we can allow that your moral obligations do not change when the sort of dependence at issue is epistemically indeterminate, even though what you are responsible for is determined by facts beyond your control.

¹⁵⁴ This is not to downplay in anyway the seriousness of the harm Ms. Genovese's assailant caused.

¹⁵⁵ All of these remarks take for granted—perhaps charitably—that *had* someone intervened on Ms. Genovese's behalf, her death would have in fact been prevented.

II. Coordination and Cooperation

Two related concepts follow in the wake of a discussion of one-to-group dependence: coordination and cooperation. When thinking about one-to-group dependence, what is the relationship between these terms? And how does this relationship cast light on some of the moral issues raised by one-to-group dependence? On the one hand, whether or not members of the group are required to cooperate or merely coordinate depend entirely on the contingent facts of the situation—or so I will argue. Such facts do not alter the underlying moral duty to rescue the victim. On the other hand, to say that the group has a moral duty to rescue the victim underdetermines who among the group ought to respond when not all the members are needed. It might be asked: it is true that *someone* in the group ought to rescue the child, but why should that someone be *me*? This question is taken up in the next section.

Cooperation and coordination, while related, are importantly distinct terms: Coordination involves bringing into proper order or relation, while cooperation consists of working with another or others for a common purpose.¹⁵⁶ It is thus possible to coordinate without cooperating and to cooperate without coordinating. In a one-to-group dependence situation, the group has the responsibility an individual would have in a case of one-to-one dependence, viz., to take the minimal steps necessary to prevent harm from coming to the imperiled individual.

So does this mean that the group has a duty to cooperate? Not necessarily. In the case of **Crowded Beach** there are multiple ways in which your group can

¹⁵⁶ Thus cooperation requires a higher level of group interaction than coordination. You and I coordinate our behaviors when each of us drives on the right-hand side of the road; this does not require working together so much as fixing on and following a shared practice or rule. To cooperate, however, is to *combine* our efforts to bring about a common goal, and this requires working together.

discharge its duty without cooperating. For example, when one beachgoer, having processed the situation more quickly than the rest, enters the surf without warning, we do not say that you cooperated in bringing about the child's rescue. But saying this does not imply that had that beachgoer not acted instinctively, none of the other beachgoers would have responded at all. The reason the rest of you do not attempt a rescue is not simply that someone is already attempting a rescue; it is that *in a situation like this*, where one person's actions are more than sufficient for rescue, your attempting a rescue at the same time could threaten the success of the rescue. Thus while it is true that you as a group share a common purpose (i.e. rescuing the child), it is misleading to say that you are necessarily working together to bring about that end. Rather, each of you is coordinating or adjusting his behavior in the appropriate way, even if it amounts to little more than restraining yourself from acting. But some situations require something more; they require cooperation, as is the case in **Frozen Lake**.

A young man has broken through the ice in the center of a frozen lake. The group has a chance of saving the young man's life, but only if the members cooperate. The number of ways in which the group can achieve its end is considerably limited; it is not enough that each member adjusts his behavior as a response to what other members do. You must work together. For example, one of you (perhaps the fastest) ought to run to the nearest house to notify emergency personnel. Suppose a coiled rope is discovered on a nearby dock. The lightest among you ought to make his way out onto the lake surface just far enough to throw the rope to the young man, while the strongest members of the group ought to prepare

themselves to haul in the victim. In a situation like **Frozen Lake**, one-to-group dependence requires more than coordination; it requires cooperation, that is, the combining of efforts to produce an end.

The moral then is not surprising: whether a situation of one-to-group dependence requires coordination or cooperation depends on the contingent facts of the situation. Morally speaking, it doesn't matter *how* you all save the child, just so long as you all save the child. Some rescue situations require only the actions of a single individual (as in **Crowded Beach**); coordination (albeit tacit) is often sufficient. Some rescue situations require multiple rescuers working in tandem; here, quite obviously, cooperation is necessary. These remarks, as it happens, are not misapplied to cases of one-to-one dependence, where someone is uniquely dependent on me for survival. Intrapersonally, I may confront conflicting instincts: there is a part of me that may dread (or resent or fear) the responsibility that unique dependence may suddenly place on me, and this response may encourage me to seek reasons to ignore my responsibility. This instinct may be countered by another, perhaps self-aggrandizing, response, one that threatens to go too far in the other direction: I may seek to dramatize my rescue in ways that threaten the safety of both me and the victim. In any event, what is required is that I bring these various impulses into harmony in whatever way will ensure that I fulfill my duty.

III. Partial Compliance

The general thesis has been that each member's responsibility to the group derives from the group's responsibility to the victim. But this rather simple

contention is of little guidance when the assistance of only some members of the group is needed. Just because the group has a responsibility to rescue the victim does not explain why I should be the one to make the needed sacrifices to rescue the victim—instead of some other member of the group. The general thesis is also unhelpful in explaining what my obligations are when others refuse to contribute, leaving me to take up the remaining slack. To answer these questions requires drawing on earlier remarks concerning determinacy: a determinate individual facing serious harm has a claim on me that is not outweighed either by the arbitrariness of my responding to him when others could or by the unfairness of my having to respond because others refuse. On the other hand, when there is no determinate individual dependent on my actions, as in the case of humanitarian relief, what I am required to for others is limited by the general level of compliance. In short, a compliance constraint along the lines suggested by Murphy (see Chapter 4) is reasonable only where there is no determinate individual who has a claim on my actions.

In a way this brings us to the center of the debate over the demands of beneficence, at least in the way it has transpired in the literature. For unlike cases of one-to-one dependence, where there is little doubt about what you are required to do (even if there are disagreements over what grounds those requirements), cases of one-to-group dependence introduce a second sort of question, one that does not arise in the former case, viz., what sort of moral relationship do I stand in to *other* (potential) donors? I am bound to the other members of the group in the minimal sense in which the group represents the only entity capable of preventing serious harm to some

individual (at little cost or risk to itself). In this sense I share with the other members a common purpose. But where fulfilling this purpose requires less than joint participation, or where joint participation is simply not viable, then the members of the group will assume an uneven balance of burdens. So what, if anything, justifies such an uneven balance of burdens?

Now there's no pressure to think that this question requires an answer in every situation, for in situations like **Crowded Beach** there may not be an answer. Is it obvious that the individual who eventually rescues the child has shouldered a burden the rest of us have not? Most likely, she will reap rewards the rest of us will not; after all, she saved someone's life at some risk to herself, and we tend to refer to such people (correctly or not) as *heroes*. This is only to suggest that the question of demands will not arise in every case of beneficence since some (many?) cases of beneficence do not incur what we typically think of as burdens. But surely some cases do incur costs—humanitarian relief is perhaps the most salient sort of case.

The concern at issue, however, need not be spelled out in terms of costs. It may be that my being saddled with the group's responsibility is *unfair*. As a rather extreme example, consider an amended version of **Crowded Beach**: four of you announce that you have no intention of rescuing the child, apparently leaving me as the only potential rescuer.¹⁵⁷ Since the child is uniquely dependent on the group, and since I am a member of the group, I have a reason to participate in an effort to rescue the child. It is true that this is a reason shared by each member of the group; thus, each has a share in the group's eventual decision. But when you announce your

¹⁵⁷ Imagine that we are avowed racists and the drowning child is a member of a despised race.

intentions to let the child drown, your shares fall to me. And this it seems is unfair, even if (let us assume) I had every intention of entering the water myself.

Suppose, instead, the case is **Restaurant**: four of you announce at the end of dinner that you have no intention of tipping the waiter. (You hold, let's say, the misguided belief that this will force an increase in service-industry wages.) Since there is no reason for refusing to tip, and since the waiter indeed deserves a tip, the burden falls on me to supply not only my share of the tip, but your shares as well. This is plainly unfair. It is unfair because I have to make up for your non-compliance. But identifying this source of unfairness does not settle the issue. Indeed, the issue only now comes into focus: should this unfairness affect what is owed to the recipient? In other words, does the unfairness done to me in anyway diminish my responsibility to the recipient?

Liam Murphy (2000), as you'll recall from Chapter 4, thinks it does. He rightly points out that requiring me to take up the slack for those individuals who do not comply places an unfair burden on me. The question, however, is whether that burden is sufficient to deprive the recipient of what she is, by all accounts, owed. Murphy is prepared to say that it is, but it is not at all clear why the *recipient* should inherit, as it were, the unfairness done to me. After all, refusing to take up the slack for the non-compliance of others does not rectify or offset the unfairness done to me; it only punishes the recipient(s) of aid.¹⁵⁸ So while it may be unfair that the burden falls to me to rescue the drowning child in **Crowded Beach** when the other

¹⁵⁸ The most vivid depiction of this (see Chapter 4, section III) involved five drowning children and five (potential) rescuers. On Murphy's account, morality requires only that you save one of the victims, *even if* the other rescuers refuse to comply. But if morality requires anything, it requires that you make an effort to rescue the other four children, so long as doing so does not seriously jeopardize your safety.

beachgoers refuse, that unfairness is not a reason to refuse to rescue the child, since the child had and continues to have a claim on the group in virtue of his being uniquely dependent on the group for his survival. Likewise in **Restaurant**: it is indeed unfair that the burden falls to me to cover the tip for the non-compliers, but this is not a reason to deprive the waiter of what he is owed.

Let me hasten to add, however, that in the event that I see your non-compliance as a reason not to tip or as a reason not to rescue the child, the responsibility for the harm does not rest solely on my shoulders; it rests on the *group*. All of us are culpable, including myself. I do not, in other words, shoulder a disproportionate share of the blame, even if I am asked to assume a greater share of the burden. Rather, since the group has a (perfect) moral duty to rescue the child, the group is jointly responsible for his subsequent death. So while each of us may see different reasons for refusing to aid the child, none of these reasons removes the group's obligation. Quite obviously, the laziness or racism that stands behind your reason to refuse to aid the child does not remove the group's obligation, but neither does my reason, viz., that I am left to take up the slack left by your refusal.

What is of critical importance here is the fact that there is an individual—the drowning child in the amended **Crowded Beach** case and the waiter in **Restaurant**—who has a compelling moral claim on the group, and none of the reasons adduced by the members of the group outweighs the reason that that individual has.¹⁵⁹ So long as none of us risks losing something of significance by aiding the victim, the fact that there is a person to whom we must answer transforms the moral equation. As I

¹⁵⁹ Admittedly, the waiter's claim may not seem to carry much moral weight, and it is perhaps odd to pair his plight with that of the drowning child. It is, however, the symmetry between his reason and the child's reason that I wish to highlight.

argued in Chapter 5, persons are the source of reasons, and respecting persons entails respecting their reasons. In the cases we've been considering, we are not weighing the burdens on me against the value of beneficence, impersonally construed; we are weighing the burdens on me against *his* reasons for wanting to be treated decently.

It would be a mistake, however, to abandon Murphy's insight altogether. By attending to the unfairness done to you by our non-compliance, we can explain the intuition that at least some principles of beneficence are unacceptable. But which ones? They will not be the ones involving determinate individuals whose claims on us are morally conclusive. They will not be ones where the risks or costs of complying do not rival the risks faced by the victim. The unacceptable principles of beneficence, I contend, are those that place substantial, unfair burdens on me as a result of the non-compliance of others, but where *there are no determinate individuals* who can claim that my refusal to take up the needed slack cost them something significant. What I have in mind here, principally, are the costs of humanitarian relief.¹⁶⁰ It is here that I believe Murphy's compliance constraint finds a natural home.

By letting determinacy carve up principles of beneficence in this way, we can explain two very robust intuitions. First, all will agree that there are situations in which the non-compliance of others can directly impact what beneficence requires of me. Just because under full compliance I would be required to rescue only one of the five children drowning in front of five of us does *not* mean that in the event that you four refuse to do your share I am not required to make a good faith effort to rescue the other children (when doing so does not seriously threaten my own well-being). No

¹⁶⁰ See Chapter 3, section 2.

one can seriously maintain this view. The second intuition my approach can explain is the one most of us have in reaction to the revisionist views of Singer and Unger. The thought is that while we ought to divert some of our income to humanitarian relief efforts, we are not obliged to center our entire lives around the single aim of eradicating extreme poverty, as important as that aim is. We are not required, for example, to live lives at or just above the poverty line. But it should be emphasized at once just *why* humanitarian relief involves such extreme sacrifices: because most individuals are not doing their fair share.

This raises the question: what sort of sacrifice would each of us have to make under full compliance? That is, in order to eradicate extreme poverty (by World Bank standards, this amounts to an individual's living on less than a dollar a day) what would each individual in a position to contribute have to give if *all* such individuals intended to give? Answers to this question of course vary, but best-guess estimates will come as a surprise.

According to the head of the United Nation's Millennium Development Goals, if the U.S., which currently spends 15¢ of every \$100 of our national income on humanitarian assistance, increases its contributions to 70¢ of every \$100 (0.7% of GNP), and if other developed countries meet the 0.7% of GNP goal, eight million unnecessary deaths could be prevented every year.¹⁶¹ The World Health Organization Commission on Macroeconomics and Health (2001), a commission involving more than a hundred specialists from around the world, calculated that donor aid ought to

¹⁶¹ See Sachs (2005) and CMH (2001). Sachs of course has his detractors, many of whom express skepticism about the details of his proposals. But even if Sachs is wrong—evenly substantially—about specific figures, there is no doubting the difference between what one must give under full compliance and what one must give under minimal compliance. And this is the issue of philosophical importance in the text.

rise from around \$6 billion per year to \$27 billion per year. “With a combined GNP of the donor countries equal to around \$25 trillion dollars as of 2001, the commission was advocating an annual investment of around *one thousandth of rich-world income*” (Sachs, 2005: 204). For the U.S., this would amount to an increase of \$60 billion a year. How could this revenue be generated? There are a number of strategies to choose from; while none is politically popular, all are practically possible: a flat increase in federal taxes, the imposition of a luxury tax, re-appropriation of the federal budget, to name just a few. Bear in mind that nearly \$4 billion of our GNP is currently spent on cotton subsidies for U.S. farmers. We currently spend half that amount on humanitarian relief. So, for example, redirecting cotton subsidies to humanitarian relief would further reduce our individual burdens, even as it increased the burdens on cotton farmers.¹⁶²

It is also worth keeping in mind that the top 400 income earners in the U.S. saved \$7 billion dollars after President Bush’s recent tax cuts were enacted.¹⁶³ While one might argue that requiring these individuals to donate their tax savings to humanitarian relief would be unreasonable—though it’s not clear how that argument would go, seeing as these individuals have amassed all the earthly possessions one could possibly want—it would hardly be unreasonable to propose rolling back (at least part of) the president’s tax cuts in an effort to meet the U.N.’s Millennium Development Goals. \$75 billion dollars is a drop in the bucket when one considers

¹⁶² A more important consequence of reducing cotton subsidies would be to spur agricultural and economic development in those eastern and central African nations that depend on cotton. See Amadou Toumani and Blaise Compaore, “Your Farm Subsidies Are Strangling Us,” Op-Ed, The New York Times July 11, 2003, A19.

¹⁶³ The combined income of these individuals totaled \$69 billion, more than the combined incomes of 166 million people across four African nations (Nigeria, Uganda, Senegal, and Botswana)

the hundreds of billions of dollars in revenue the government would have collected under the previous tax system.

It is important, however, not to lose sight of the larger philosophical point: Whether or not our figures are trustworthy, what each of us would be required to do to effectively eradicate extreme poverty under full compliance is but a tiny fraction of what we would have to do to eradicate extreme poverty under minimal compliance. After all, under minimal compliance, we would be required to live lives at or just above the poverty line. Even if rich countries needed to contribute 1.7% of their GNPs, *twice* the current estimate, this would hardly compare to the kind of burden each of us would shoulder under partial compliance. There is then a strong case to be made that my being required to live a life at or just above the poverty line is unfair when the *reason* for my living such a life has everything to do with the fact that others are not doing what they should. This is not simply covering everyone else's tip after an expensive meal; this is spending the rest of one's life washing dishes to make up for the difference. Surely there is a lesson in this. But what is it?

One might be tempted to raise the same concern that dogged Murphy's proposal: does the unfairness done to me constitute grounds for depriving individuals of much needed assistance—who, by the way, have not acted unfairly toward me? We might reasonably wonder why individuals abroad should be deprived of life-saving assistance just because I am treated in a way that is objectionable. Indeed, this seems to represent yet another instance of the fallacy attributed to Murphy. From the fact that others do not treat me fairly insofar as I am left to take up the slack of their

noncompliance, it does not seem to follow that the imperiled individual is any less entitled to my assistance when she stands to suffer a serious loss.

The reason this does *not* represent yet another instance of the fallacy attributed to Murphy has to do with the fact that in the case of humanitarian relief, there is no determinate individual who is entitled to *my assisting him*. Instead, as I argue in the next section, for any potential recipient of humanitarian assistance, there are two reasonable objections she might have to my not contributing to humanitarian relief. She can claim, first, that my not contributing lessens her chances of survival by some fraction, but this leaves open the possibility that I may have reasonable grounds for allowing her to assume greater risk.¹⁶⁴ Second, she can claim that her survival depends exclusively on the actions of wealthy nations of whom I am a member. This latter claim amounts to the claim that I fail to contribute to the group's humanitarian responsibility, which implies that I do not discharge my share, if any, of the group's responsibility. What is essential for the purposes of understanding the difference between what we might call humanitarianism and Samaritanism, is that the first claim—but not the latter—does not necessarily block any objection that I might have to a principle requiring that I take up the slack for others. I may have, in other words, an objection that it would be unreasonable for any potential recipient of humanitarian aid to ignore.

The difference can be seen by returning to several cases we considered in Chapter 3. In the next section, I attempt to bring out this difference in the context of Scanlon's framework for moral justification, suggesting further how this might aid us in formulating an acceptable principle of humanitarianism.

¹⁶⁴ See Chapter 3.

IV. Determinacy and Scanlonian Contractualism

Draw your attention back to the hypothetical cases introduced in Chapter 3:

Detention Center I. Terrorists are about to strike a youth detention center where 2,000 youths are detained, each in a separate cell. As it happens, cell doors can only be opened electronically, from a centralized computer. As part of the attack, however, the terrorists have disabled the computer in such a way that only 999 cell doors will be opened automatically before the center's destruction. For unknown reasons, however, the computer will select the 999 doors *at random*. So until the designated moment at which the randomizer selects which doors to open, there is no fact of the matter as to which doors will open and, hence, which youths will be saved. There is, however, something you can do in the little time remaining: you can enter into the computer a lengthy code that will instruct the system to open 1,000 cell doors, instead of 999—though, again, which doors will open will be the result of the randomizer.

Detention Center II went like this:

Next to the computer that controls the 2,000 cell doors is a second computer that controls a single cell door, separate from the rest, and which is unaffected by the operations of the main computer. There is no indeterminacy about who will be evacuated if you decide to enter a different (though equally lengthy) code into this second computer, viz., the occupant of cell #2000, say.¹⁶⁵ Assume as before that even without your assistance 999 cell doors in the main detention center will open

¹⁶⁵ To keep the numbers straight, imagine that the main computer now controls 1,999 cell doors, instead of 2,000. In both examples, then, there are exactly 2,000 individuals whose lives are at stake. Of course this change is not entirely without consequence: now, for any youth in the main detention area, the probability of his being saved without your assistance is now 0.4998; with your assistance, 0.5002.

automatically. Now imagine that you only have enough time to enter one code into one computer. What should you do? On Option 1, you enter the code into the main computer. On Option 2, you enter the code into the second computer.

From a completely impersonal perspective, that is, without regard to how your decision would affect any particular individual, Option 1 and Option 2 appear to be equally permissible. On both options the result is 1,000 lives saved and 1,000 lives lost. But when we move to the perspectives of the individual detainees, this equivalence disappears. For from the perspective of any one detainee in the main detention center, his chances of survival if we choose Option 1 are four-hundredths of one percent better than his chances of survival if we choose Option 2. From the perspective of the occupant of cell #2000, however, the difference between Option 2 and Option 1 is the difference between life and death.

I argued at length in Chapter 3 that we should regard #2000's objection to Option 1 as a legitimate objection, one that ought to influence our decision. Whether it is you or me or a total stranger occupying cell #2000, that person would have that reason, simply in virtue of the circumstances, and that reason would form (counterfactually, if you like) the basis of a compelling objection to a policy allowing Option 1 over Option 2. I will assume, then, that we ought to hold that that consideration is more than enough to break the tie between Options 1 and 2. The more general lesson seems to be that moral reasoning is sensitive to the sorts of considerations that arise from the particular individuals involved. But such a sensitivity calls out for explanation.

Scanlon offers an appealing suggestion: determinacy is entailed by a basic constraint on moral justification, viz., that it is only individuals that can accept or reject principles. “[T]he guiding idea of contractualism,” writes Scanlon (1998: 229), and “what enables [contractualism] to provide a clear alternative to utilitarianism and other forms of consequentialism” is the idea that “the justifiability of a moral principle depends only on various *individuals*’ reasons for objecting to that principle and alternatives to it.” The main motivation behind this idea is the concern that aggregative benefits could outweigh the costs imposed on an individual or minority, a concern that is often associated with utilitarianism and various forms of consequentialism. For this reason, Scanlon (1998: 230) insists that “all objections to a principle must be raised by individuals.” Note that this constraint delivers the intuitive results we’ve been considering.

This explains, for example, why we ought to choose Option 2 over Option 1 in **Detention Center II**: #2000’s objection to a principle allowing us to choose Option 1 is far stronger than any one detainee’s objection to a principle allowing us to choose Option 2. Mutatus mutandis for **Detention Centers III** and **IV**. But this also can explain why we ought to enter the code in **Detention Center I**. I do not have a reasonable objection to a principle requiring me to enter the code, since there is no risk or cost to me of entering the code, whereas any one detainee *does* have a reasonable objection to a principle allowing me to refuse, viz., that refusing decreases his chances of survival, even if only minimally. Let us then apply some of these considerations to the question of humanitarian relief.

In considering the moral implications of various principles requiring humanitarian assistance, the point to bear in mind is that all objections to a principle must be raised by individuals. This immediately places constraints on how principles are rejected: principles cannot be rejected on the basis of the *sum* of benefits or burdens to a group of individuals.¹⁶⁶

Consider then a principle allowing me to spend money entertaining friends Labor Day weekend, money that could otherwise be contributed to humanitarian relief efforts.¹⁶⁷ Which *individual(s)* could reasonably reject this principle? Our untutored response is: those individuals whose lives would be saved by that contribution. But as I've argued, there are no such individuals—just as there is no determinate individual whose life would be saved by entering the code in **Detention Center I**.

Now this does not mean that an individual living in absolute poverty (*A*) has *no* objection against my spending money entertaining friends Labor Day weekend.

¹⁶⁶ Scanlon goes so far as to assert that “the sum of the smaller benefits to others has no justificatory weight, since there is no individual who enjoys these benefits and would have to forgo them if the policy were disallowed” (ibid., p. 230). This claim, however, invites the following objection:

Detention Center V: You can either save 1,000 individuals or you can save the occupant of cell #2000, but not both. If you like, assume that the 1,000 individuals are determinate, such that there is a fact of the matter as to who will die if you decide to rescue #2000 (Option 2).

Intuitively, one ought to save the 1,000 individuals. Trouble is, it's not clear how, or if, someone like Scanlon can arrive at that judgment, since the sum of benefits (or harms) has no justificatory weight. If we consider the matter simply from the perspective of the particular individuals involved--#2000's objection to Option 1 versus anyone else's objection to Option 2—the two options appear morally equivalent; we are, after all, prohibited from comparing the sum of the consequences. Scanlon attempts to meet the challenge by identifying a consideration that any one in the main detention center has that the occupant of cell #2000 does not: Option 2 does not take into account the value of saving his life, “since it permits [you] to decide what to do in the very same way that it would have permitted had he not been present at all, and there was only one person in each group” (p. 232 ff.). And on these grounds, Option 2 is reasonably rejectable. This response is not entirely satisfactory. It is, however, a coherent response, and one that will have to suffice for now.

¹⁶⁷ As I argue below, this principle might be rejected simply on the grounds that there is an alternative principle (e.g. principles of taxation) that would be both more effective as well as fair. Here I am exploring principles favored by revisionists like Singer and Unger.

But the content of *A*'s objection is analogous to the content of the objection any one detainee could raise against your not entering the code in **Detention Center I**: by not taking the relevant action, I decrease his chances of survival by some fraction. While this fraction is of course non-negligible, it does not make comparison with other pursuits—for example, how I might choose to spend my money on a given weekend—illegitimate. It is an open question, at worst, whether or not *A* could reasonably reject a principle allowing me to spend money entertaining friends Labor Day weekend when it is unlikely in the extreme that that money would have any direct impact on *A*'s condition, given contingent facts about the structure of humanitarian relief organizations. To take this result seriously is to conclude that, as matters stand with respect to humanitarian relief efforts, we could reasonably reject a principle of humanitarianism that required that each of us contribute on every available occasion. Because *A*'s complaint against my not contributing amounts to an objection to my not decreasing the overall amount of suffering in the world, of which *A* is indeed a part, there would seem to be many pursuits and goods (but certainly not all) that this complaint cannot reasonably prohibit.

My objection to a principle prohibiting me from entertaining friends Labor Day weekend is stronger than *A*'s objection to a principle allowing me to entertain friends that weekend. Why? Because the difference between my conforming to the former principle and my conforming to the latter principle, *as far as A is concerned*, is negligible. It may improve the overall state of affairs, but this is a separate matter from improving *A*'s condition. But the difference between my conforming to the former principle and my conforming to the latter principle, as far as I'm concerned, is

not negligible, since entertaining friends is, arguably, a pursuit whose value contributes meaningfully to a life worth living. Matters of course would be different if there was a determinate individual whose life hung in the balance, but again this is not the case before us.

Now some may not find sacrificing time with friends an imposition, but this simply means that conforming to a principle prohibiting entertaining friends is supererogatory since it is enough that *A* cannot reasonably reject a principle allowing these token acts. Thus I conclude that with respect to a principle of humanitarianism requiring contributions to humanitarian relief efforts *on every available occasion*, the distinction between determinate and indeterminate beneficiaries, together with a Scanlonian model of moral reasoning, yields the plausible conclusion that such a principle should be rejected. And notice that this rejection does rest principally on the demandingness of the various principles.

Consider, however, a principle that never requires affluent countries, considered as a group, to contribute to humanitarian relief. On this principle, groups are allowed to spend their disposable incomes on any pursuits and goods they please whether or not it includes humanitarian aid. What sort of objection, if any, could an individual raise against this principle? To answer this question we must first consider how alternative principles might affect *A*'s condition and what the difference *to A* is between such principles. So imagine a principle requiring that every group give a modest amount of its income to humanitarian relief. Such an amount would not require a group's members forgoing important projects and would even allow for occasional luxuries. The question then is this: would the difference between their

conforming to the former principle and their conforming to the latter principle be negligible, as far as *A* is concerned? Now I am supposing that, as far as the members of the groups are concerned, the difference between these two principles is negligible.¹⁶⁸ But as far as *A* is concerned, it's almost certain that the difference is not negligible (see Appendix). For when we consider how critical the services to poverty-stricken areas are, and how these services can only be provided by large donor organizations (including nations), it's very likely that there are individuals whose conditions were improved as a direct result of these group contributions, such that had those groups not made those contributions, there is a considerable chance that their lives would have been worse.¹⁶⁹ If this is indeed the case, then in the absence of some story to the contrary, a principle of humanitarianism requiring that groups make modest contributions to humanitarian relief could not be reasonably rejected. In effect, this is what the U.N.'s Millennium Development Goals require: 70¢ of every \$100 a nation earns. The upshot, however, is that while I can reasonably reject many principles requiring this or that *token* contribution, I cannot reasonably reject a principle requiring groups to contribute to humanitarian relief.

This, however, leaves open the following question: in what sense am I required to participate in such a group? In other words, can I reasonably reject a principle that required that I participate in a group among whose moral mandates include humanitarian relief? To answer this question requires more than simply noting such generic features as my rational agency, my conception of the good, etc. It requires taking account of the many features specific to my life that may make a

¹⁶⁸ For example, CARE will automatically deduct a specified amount from one's bank account each month, so the process of contributing is as non-intrusive and effortless as a sales tax.

¹⁶⁹ See Sachs (2005), Chapter 15.

principle requiring participation unreasonable. As it happens, these features do not make such a principle unreasonable. Indeed, I suspect this is the case with most people in affluent countries. But the procedure outlined here does not presume that every individual is under the same moral obligation. What it does assume is that a principle permitting affluent groups to refuse to contribute to humanitarian relief is reasonably rejectable. As a parallel, consider **Crowded Beach**: there is no moral excuse for the *group* to refuse to attempt to rescue the child, even though there may be reasons why *you* should not have to participate here and now.

With respect to humanitarian relief, the story is familiar. You, qua private citizen, do not have a perfect duty to any one individual abroad who may be suffering unnecessarily since there is no determinate individual whose life depends on your efforts. At the same time, when we conceive of ourselves as a group (as a nation, in this case), this group does have a perfect duty to participate in humanitarian relief efforts. The question of interest then becomes: how should we understand our relationship to this entity, and what sort of duties, if any, do we have vis-à-vis this entity? I have argued (Section II) that our obligations will vary according to the contingent features of the situation. That we as a group have a perfect duty to respond to the extreme poverty abroad does nothing to settle *how* we should respond, as any familiar with current development economics debates can attest to.

There is a growing consensus among economists and some philosophers that the only realistic means of addressing (let alone, eradicating) extreme poverty lie at the institutional level.¹⁷⁰ The forces that shape the conditions in the world's most destitute places (principally, Africa) are institutional forces. This is not just because

¹⁷⁰ See Pogge (2002).

institutions like the World Bank and International Monetary Fund command large sources of capital; it is also because they are responsible for setting policy—lending policies, debt-relief policies, trade policies, and so on. According to many, how the rules of the game are set down makes all the difference. For example, a nation's ability to improve and maintain fragile infrastructures depends critically on how much money that nation is able to borrow and the conditions under which it must repay that money. Even more critical in the minds of some economists are the trading policies set forth by leading nations, for whether or not a developing nation can trade on international markets (and under what restrictions) largely determines how much capital, if any, a country can accumulate. And the ability of a country to accumulate capital translates almost immediately into more jobs, better health services, better education, longer lifespans, and so on.

Then there are the institutional forces of independent nations, particularly the wealthy nations of the West. So, for example, how the money in our federal budget is distributed among its various agencies and services can make sizable differences in the living conditions of those abroad: the nearly \$4 billion in cotton subsidies are a prime example. But distribution is only part of the story, for whether a greater proportion of the budget goes to cotton subsidies or humanitarian relief projects, there is the more basic question of how much money the government has at its disposal. And, as we all know, how much money the government has at its disposal is largely (though not entirely) a function of tax revenues. So federal spending on USAID (for example) is not simply a matter of its percentage of the federal budget; it's also a matter of how big the budget is. It is not simply a question of how resources ought to

be distributed; there are very important—arguably, more important—questions about domestic tax structures, international lending policies, debt relief, and so on, that need to be addressed.

V. Good Humanitarians and Good Samaritans

Let me close with some general remarks concerning the distinction between humanitarians and Good Samaritans, for in one sense this has been one of the overarching aims of this work, viz., to limn the crucial differences between our obligations to those (determinate) individuals in emergency situations and those living in extreme poverty. Being a Good Samaritan requires responding appropriately to the needs of those uniquely dependent on you, for you have a conclusive reason to assist someone who is uniquely dependent on you, if you do, because it is partly constitutive of our moral attitude towards others that we see a reason not to view their lives as open to exchange. Part of this involves the fact that there is a determinate individual to whom we must justify ourselves. In the drowning child case, there is an individual to whom you must answer. And any attempt to dismiss his expectation that you will take the minimal steps necessary to prevent his death, *when no other individual can be expected to do so*, is unreasonable in the extreme. **Vintage Boots** nicely illustrated this point.

But the intermediation of humanitarian relief agencies changes our moral relation to distant sufferers. Humanitarians, unlike Good Samaritans, are not responding to the needs of any particular individuals; the structure of humanitarian relief organizations prevents it. Rather, they are responding to the state of humanity,

to the condition of “a needy people.” One cannot take responsibility for any particular individual abroad, short of intervening directly. One is simply not in control of any distant individual’s life. We have little choice but to view the work of humanitarian relief organizations as directed at the *people* whose lives are most at risk. This, it seems to me, is a natural gloss on “humanitarian.” And when you make no effort to feed the humanitarian relief machine, you ignore a reason to improve what can only be described as a needy people; this is quite distinct from failing to see a reason to assist someone whose death is virtually assured if you do nothing to prevent it.

The notion of betraying humanity is importantly distinct from the notion of betraying individuals uniquely dependent on us. To *betray* is to break faith with, to fail to meet the hopes of. What I have attempted to establish is that failing to meet the hopes of a distant people is importantly different than failing to meet the hopes of an individual whose life is, quite literally, in your hands.

APPENDIX

A Universe Apart

IT HAS BEEN SAID that we in the First World represent the “Golden Fraction,” for while the majority of the world’s population live in poverty (much of it extreme), a slim fraction of us, roughly 15%, enjoys living standards of relative luxury. To say that our quality of life is “a universe apart” from the quality of life most of the world’s inhabitants are used to is hardly an overstatement. The following numbers tell the story.

First, the more familiar facts—those readily available from the World Health Organization: Nearly a third of the world’s population is infected with tuberculosis. About two million of those individuals die every year of the disease. To put this into perspective, consider that 200 people die of tuberculosis every *hour*. Malaria claims more than one million lives a year; that’s more than 100 people an hour. Diarrheal diseases such as cholera, dysentery, and E. coli, result in about 1.9 million deaths a year, most of those infants and young children. (This is hardly surprising given that 1.5 billion people do not have access to clean water.) Nearly 900,000 individuals a year die from measles. The hepatitis B virus claims one million lives a year.

It should go without saying that we in the First World haven’t seen numbers like these in decades, if not centuries.¹⁷¹ And there’s a good reason for this: in the United States and other wealthy countries, annual public spending on health care

¹⁷¹ The Center for Disease Control identifies the following diseases as causes of death in the US in 2000 (number of deaths are given in parentheses): influenza (1,765), TB (776), malaria (3), measles (1).

averages about \$2,000 per person. In Africa, the average is around \$10 per person.¹⁷² Contrast as well the differences in access to physicians: according to a World Bank study, in the industrialized countries of the West, there is one physician for every 450 individuals; in Sub-Saharan Africa, there is one physician for every 23,850 individuals. It is no surprise, then, that individuals in Sub-Saharan Africa can expect to live nearly *thirty-five years less* than individuals in the United States and Britain. These individuals live what Ted Honderich (2003) calls “half-lives.”

But it is the AIDS virus that promises the greatest devastation. An estimated 42 million people are H.I.V.-positive. In China, Ethiopia, India, Nigeria, and Russia, the number of AIDS cases is expected to double by 2010, with a total of 50 million to 75 million infected people in *those countries alone*.¹⁷³ Last year the virus claimed more than 3.1 million lives, and the number is only expected to rise. In Africa, 29 million individuals are living with H.I.V. or AIDS.¹⁷⁴ In some African countries the rate of infection is as high as forty percent. But no one can fully anticipate the scope of the ruin the virus is likely to bring to the African continent, for when we consider the already precarious conditions under which so many endure, it becomes hard to wrap one’s mind around the coming holocaust. For it is estimated that, in the next seven years, 25 million African children will be orphaned as a result of AIDS. (That exceeds by several million the population of New York City.) And in a place where

¹⁷² Sachs (2003). A New York Times editorial (July 28, 2003) reports that annual health care spending in Ethiopia averages \$1.50 per person.

¹⁷³ Kramer (2003).

¹⁷⁴ Kristof (May 20, 2003).

malnutrition is already rampant (50 million children under five), the rate of starvation the AIDS virus will precipitate is expected to skyrocket.¹⁷⁵

Quite apart from the ravages of disease, however, there is chronic malnutrition. According to the United Nations Development Program (UNDP) 800 million people in the Third World do not have enough to eat. There is as well the perpetual threat of famine. For example, as of this writing, 40 million East Africans are at risk of starvation. Nicholas D. Kristof, who recently traveled to East Africa to report on the coming famine, reported that “in the best of circumstances, about 100,000 boys and girls...will die of malnutrition-related ailments this year in Ethiopia.”¹⁷⁶ And famines tend to follow on the footsteps of mass migrations. According to a recent report by the US Committee for refugees, there were 13 million refugees or internally displaced persons in Africa at the end of 2000.¹⁷⁷

As staggering as they figures are, we must not allow them to eclipse the more fundamental disparity between the First and Third World. And that is, not surprisingly, economic. The dismal health conditions among the world’s poor are just a symptom of ruinous economies. I have often wondered why so few revisionists include a separate discussion of the economic disparity between rich and poor, for it is here, if no where else, that one struggles to understand how we can continue to ignore the plight of literally billions of individuals. Perhaps it is because, as the former Nicaraguan president Daniel Ortega put it, “the field of economics is very complex and very boring.” Unlike widespread starvation, there is nothing

¹⁷⁵ See www.unicef.org

¹⁷⁶ Kristof (May 23, 2003).

¹⁷⁷ U.S. Committee for Refugees (2001), p. 15.

particularly provocative about per capita income. In any event, consider the following figures.

More than half of the world's population (approximately 3.1 billion people in 1990) had an average annual income in 1990 of \$350 US dollars.¹⁷⁸ In the poorest countries of the world (e.g. Burkino Faso, Chad, Somalia, Sudan, Bangladesh, Laos, Nepal) per capita incomes hovered below \$250—about 69 cents a day. In India, 55 percent of the population lives on \$370 a year. In Sub-Saharan Africa, the percentage is 47. By contrast, the average annual income of a U.S. citizen in 1990 was \$21,790. So just over ten years ago, the average annual income for more than half of the world's population was less than *one-sixtieth* of the typical U.S. income. But the inequity is all the more striking when one considers how wealth is distributed among rich and poor. In 1985, the poorest fifth of the world's population received 1.6 percent of the world's income, while the richest fifth received 74.2 percent. Together the fourth and fifth richest segments of the population accounted for 92.7 percent of the world's income. More recently, it was reported that 400 superrich Americans currently “command more income than 166 million people in Africa.”¹⁷⁹ Keep in mind that this is more than half of the entire US population.

Of course it would be correct to point out that incomes are not directly comparable between countries. Fluctuating rates of exchange on international currency markets and differences in costs of living are just some of the factors that can exaggerate the disparity between rich and poor. All the same, even conservative estimates put our average level of income at twenty times that of the average level of

¹⁷⁸ Unless otherwise noted, all figures are reproduced from Sachs (2005).

¹⁷⁹ Sachs (2003).

income in most places in the world. When a female lace worker in the Narsapour region of India must work eight hours a day for two to three months to earn enough money to purchase a sari costing 70 rupees (about \$1.40), the details are beside the point.¹⁸⁰

I have bothered to recite these figures for two reasons. First, philosophical discussions about distant suffering must not lose sight of just what distant suffering looks like. The first step in any discussion of the subject should be a thorough rendering of just how our lives look in relation to the lives of the world's poor, for it is *our lives* and *our choices* at which the moral question is aimed. Too often philosophers overlook this. They barrel into the discussion with exotic thought experiments and principles and all manner of theoretical apparatus, leaving the reader to wonder what any of this has to do with *her*. Moreover, unlike the discussion of drowning children or injured bicyclists, the facts describing distant suffering do not call on us to respond to how things *could be* but aren't; they call on us to respond to how things *are*.

And as they stand, things are alarmingly out of balance. The second reason for going into these details is that, at least at the level of common sense, disparities of this magnitude are morally unacceptable. There is no obvious justification for the haves to be a universe apart from the have-nots. There is no obvious justification for mothers and fathers in India having to *sell* one of their children in order to feed their remaining children when the Museum of Modern Art in New York City is spending

¹⁸⁰ Bisilliat (1987).

\$858 million just to *expand*.¹⁸¹ It is one thing for a small fraction of the world's population to enjoy decent living standards; it is quite another for that same fraction to command a level of income roughly sixty times that of the rest of the world, *especially* when that disparity directly (or even indirectly) contributes to the suffering of billions of people.¹⁸² Notwithstanding one's knee-jerk cynicism, the popular bumper sticker enjoining us to 'Live simply so that others may simply live' is, for the world's desperately poor, deadly serious.

But of course we are discriminating consumers. No one wants to throw money at a problem without some hope that it will make a difference. So, naturally, if there's a chance our contributions (assuming it's monetary contributions we're talking about) won't make a difference, then the mere fact of extreme disparities is not enough to open our checkbook. No matter how moved we might be by distant suffering, so long as *there's a chance* our check won't clear the bureaucratic divide, so long as *there's a chance* it won't be transformed into bags of rice or vaccines or O.R.T. packets, we cannot legitimately be blamed for failing to contribute to relief agencies like UNICEF or Oxfam. I am surprised how often this point is raised in conversation, as if speakers were simply passing on a bit of common knowledge, as if everyone already knew how his donations would be used or (more importantly) how past contributions have actually affected the lives of the world's poor. With the exception of a developmental economist, no one who has expressed doubts to me about the effectiveness of these agencies is at all familiar with the facts.

¹⁸¹ Both reported in The New York Times: June 23, 2003 (A3) and May 6, 2003 (B1), respectively. Incidentally, the family interviewed for the former article received about 1100 rupees for their child—about \$20.

¹⁸² That the causes of world poverty lie at the feet of developed nations remains a central topic of debate within macroeconomic theory. See Isbister (1993).

During the years 1990-2000, Oxfam spent about \$1.8 billion—more than half of which came from private donations—on relief efforts, including such things as community development, emergency relief, food aid, and program development.¹⁸³ UNICEF, by far the largest humanitarian agency, had a charitable expenditure of more than \$10 billion for the years 1990-2000.¹⁸⁴ How much bang for the buck did these agencies get? Arguably, quite a lot. The 2001 United Nations Report of the Secretary-General, “We the Children: End-decade review of the follow-up to the World Summit for Children” reports substantial improvements in the living conditions of developing countries during the 1990-2000 decade.¹⁸⁵ For example, while measles continues to plague large portions of the globe, worldwide reported measles incidence declined by nearly *two thirds*. Neonatal tetanus declined by 50 percent. According to World Health Organization estimates, deaths due to diarrhoea were reduced by 50 percent. Malnutrition in developing countries declined by 17 percent. 816 million additional people obtained access to improved water supplies over the decade despite the 1.5 billion people who are still without clean water. 747 million additional people utilized improved sanitation facilities even though 2.4 billion people—including half of all Asians—lack such access. While it would no doubt be impossible to designate which projects or policies (not to mention which agencies) bear the greatest responsibility for these improvements, there can be little question that the efforts of these agencies, funded in part by the donations of

¹⁸³ See www.oxfam.org. Annual reports documenting incomes and expenditures are all available for download.

¹⁸⁴ Because the largest donors to UNICEF are governments and intergovernmental organizations (about 60%), its budget tends to be greater than Oxfam’s, which receives most of its donations from non-governmental/private sector organizations and individuals.

¹⁸⁵ United Nations (2001).

individuals, saved lives. This is the conclusion Robert Cassen (1994) and his associates come to in Does Aid Work?, widely considered to be the standard reference on the effectiveness of international aid programs.

Another misconception about agencies like UNICEF and Oxfam is that they are only involved in providing food aid and medicine, which, if it were true, would bolster the argument that they are doing nothing (or worse) about the underlying conditions hampering real development. This is not the case, however. In northern Mali, for example, Oxfam has been giving out micro-loans to establish income-generating schemes, setting up literacy and education programs, and enabling people to participate effectively in the new decentralized political system. But Oxfam's most ambitious project promises even greater returns: *Make Trade Fair* is an international cooperative aimed at reforming unfair international trade practices. By lobbying world leaders, transnational companies, and international bodies like the World Trade Organization, the agency hopes to remove the most damaging obstacle to economic development in the Third World: closed markets. Most economists agree that without some foothold in world markets, the local economy—and with it, a person's living standards—will, at best, remain stagnant.

Oxfam's effort to break open markets would seem to be on the mark judging, for example, from a recent piece by the presidents of Mali and Burkino Faso who argue that Africa's hope of becoming a "full member of the world community" depends vitally on its ability to compete on the world market. And this ability is "is seriously threatened by agricultural subsidies granted by rich countries" to their

farmers.¹⁸⁶ UNICEF, for its part, directs part of its efforts at establishing universal access to basic education, reducing gender disparities in education, and improving skills-based learning programs. And these efforts appear to be paying off: net primary school enrollment has increased in all regions, adult illiteracy has dropped, and new partnerships are being established among education providers, industry and community leaders.¹⁸⁷

¹⁸⁶ Toumani (2003).

¹⁸⁷ United Nations, *op. cit.*

BIBLIOGRAPHY

- Benditt, Theodore M. "The Demands of Justice," Ethics 95 (January 1985): 224-32.
- Bisilliat, Jean and Michele Fieloux, Women of the Third World: Work and Daily Life (Cranberry, N.J.: Associated University Presses: 1987).
- Brink, David O. "Utilitarian Morality and the Personal Point of View," Journal of Philosophy Vol. 88, No. 8 (August, 1986): 417-38.
- Cassen, Robert and associates. Does Aid Work? Report to an Intergovernmental Task Force 2nd Edition (New York: Clarendon Press of Oxford University, 1994).
- Chatterjee, Deen K. (ed.) The Ethics of Assistance: Morality and the Distant Needy (Cambridge: Cambridge University Press, 2004).
- Commission on Macroeconomics and Health, Macroeconomics and Health: Investing in Health for Economic Development (Geneva: World Health Organization, 2001).
- Cottingham, John. "Partiality, Favouritism, and Morality," Philosophical Quarterly vol. 36, issue 144 (July 1986): 357-73.
- Cullity, Garrett. "International Aid and the Scope of Kindness," Ethics 105 (1994): 99-127.
- Dancy, Jonathan. Moral Reasons (Oxford: Blackwell Publishers, 1993).
- Darwall, Stephen. Impartial Reason (Ithaca, NY: Cornell University Press, 1983).
- . "Respect and the Second-Person Standpoint," Proceedings and Addresses of the American Philosophical Association, Vol. 78: 2 (Nov. 2004).
- Dreze, Jean and Amaryta Sen. Hunger and Public Action (Oxford: Oxford University Press: 1989).
- . The Political Economy of Hunger (New York: Oxford University Press, 1990).
- Feinberg, Joel. Harm to Others: The Moral Limits of the Criminal Law (Oxford: Oxford University Press, 1984).
- Feldman, Fred. "Comments on Living High and Letting Die," Philosophy and Phenomenological Research Vol. LIX, No. 1 (March 1999):195-201.

- Fischer, John M. and Mark Ravizza (eds.). Ethics: Problems and Principles (New York: Harcourt brace Jovanovich, 1992).
- Foot, Phillipa. "Utilitarianism and the Virtues," reprinted in Scheffler (1988).
- . Natural Goodness (Oxford: Oxford University Press, 2001).
- Friedman, Marilyn. "The Impracticality of Impartiality," Journal of Philosophy No. 86 (November, 1989): 645-56.
- Gomberg, Paul. "The Fallacy of Philanthropy," Canadian Journal of Philosophy Vol. 32, No. 1, (March 2002): 29-66.
- Hampton, Jean. "Selflessness and Loss of Self," Social Philosophy & Policy, vol. 10, no. 1 (1993): 135-65.
- Hare, R. M. The Language of Morals, Oxford: Oxford University Press, 1964.
- . Moral Thinking: It Levels, Method, and Point (Oxford: Oxford University Press, 1981).
- Herman, Barbara. "The Scope of Moral Requirement," Philosophy & Public Affairs Vol. 30, No. 3 (Summer, 2002).
- Hershenov, David B. "A Puzzle About the Demands of Morality," Philosophical Studies Vol. 107, no. 3 (February 2003): 275-290.
- Heuer, Ulrike. "Raz on Values and Reasons," in Wallace, et. al. (2004).
- Hill, Thomas E. Dignity and Practical Reason (Ithaca: Cornell University Press, 1992).
- Honderich, Ted. "After the Terror: A Book and Further Thoughts," The Journal of Ethics Vol. 7, No. 2 (2003): 161-81.
- Hooker, Brad. "Sacrificing for the Good of Strangers—Repeatedly," Philosophy and Phenomenological Research Vol. LIX, No. 1 (March 1999): 177-181.
- Hume, David. A Treatise of Human Nature eds. Selby-Bigge & P.H. Nidditch 2nd edition (Oxford: Oxford University Press, 1978).
- Hurley, Paul. "Fairness and Beneficence," Ethics 113 (July 2003): 841-864.
- Isbister, John. Promises Not Kept: The Betrayal of Social Change in the Third World 2nd edition (Hartford, CT.: Kumarian Press, 1993).

- Jamieson, Dale. "Duties to the Distant: Aid, Assistance, and Intervention in the Developing World," forthcoming in The Journal of Ethics.
- Kagan, Shelley. The Limits of Morality (Oxford: Oxford University Press, 1989).
- . "Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation," Philosophy and Public Affairs, Vol. 13, No. 5 (Summer, 1984): 239-54.
- Kamm, F. M. Mortality and Morality Vol. II. (Oxford: Oxford University Press: 1996).
- Kant, Immanuel. Groundwork for the Metaphysics of Moral, trans. and ed. By Mary J. Gregor (Cambridge, Cambridge University Press, 1996).
- . Metaphysics of Morals, translated and edited by Mary J. Gregor (Cambridge, Cambridge University Press, 1996).
- Korsgaard, Christine. "Two Distinctions in Goodness," The Philosophical Review, XCII, No. 2 (April, 1983): 169-195.
- . "The Reasons We Can Share: An Attack on the Distinction between Agent-Relative and Agent-Neutral Values," Social Philosophy and Policy vol. 10, no. 1 (1993): 24-51.
- Kramer, Larry. "The Plague We Can't Escape," The New York Times, Op-Ed, March 15, 2003.
- Kristof, Nicolas D. "Helping Binyam, When His Mother Won't," The New York Times Op-Ed, May 20, 2003.
- . "The Shape of Hunger," The New York Times Op-Ed, May 23, 2003, A27.
- Mackie, J. L. Ethics: Inventing Right and Wrong (Harmondsworth: Penguin Books, 1977).
- Mayerfeld, Jamie. Suffering and Moral Responsibility (Oxford: Oxford University Press: 1998).
- McGuire, Martin C. "The Calculus of Moral Obligation," Ethics 95 (January 1985): 199-223.
- Mill, John Stuart. (1863), Utilitarianism 2nd Edition (Indianapolis, Hackett Publishing, 2001).

- Murphy, Liam. Moral Demands in Nonideal Theory (Oxford: Oxford University Press, 2000).
- Myers, Robert. Self-Governance and Cooperation (Oxford: Oxford University Press, 1999).
- O'Neill, Onora. The Faces of Hunger: An Essay on Poverty, Justice and Development (London: Allen & Unwin, 1986).
- Otteson, James R. "Limits on Our Obligation to Give," Public Affairs Quarterly Vol. 14, No. 3 (July 2000): 183-203.
- Nagel, Thomas. Nagel, The View from Nowhere (New York: Oxford University Press, 1986).
- . Equality and Partiality (New York: Oxford University Press, 1991).
- Nozick, Robert. Anarchy, Utopia, and State (New York: Basic Books, 1974).
- Parfit, Derek. Reasons and Persons (Oxford: Oxford University Press, 1984).
- Pettit, Phillip and Michael Smith, "The Truth in Deontology" in Wallace, et. al. (2004).
- Pogge, Thomas. World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms (Cambridge: Polity Press, 2002).
- Rachels, James. "Active and Passive Euthanasia," reprinted in Fischer and Ravizza (1992).
- Railton, Peter. "Alienation, Consequentialism, and the Demands of Morality," Philosophy and Public Affairs 13 (1984): 134-171.
- Raz, Joseph. The Morality of Freedom (Oxford: Oxford University Press, 1986).
- . "On the Moral Point of View," in Schneewind (1996).
- . Engaging Reason: On the Theory and Value of Action (Oxford: Oxford University Press, 1999).
- . The Practice of Value (Oxford: Oxford University Press, 2003).
- Sachs, Jeffrey D. "A Rich Nation, A Poor Continent," The New York Times, Op-Ed, July 9, 2003.

- . The End of Poverty: Economic Possibilities for Our Time (New York: The Penguin Press, 2005).
- Sartorio, Carolina. “How to be Responsible for Something without Causing it,” Philosophical Perspectives 18, Ethics, 2004.
- Scanlon, T. M. What We Owe to Each Other (Cambridge, Mass.: Harvard University Press, 1998).
- Scheffler, Samuel. The Rejection of Consequentialism (Oxford: Oxford University Press, 1982).
- (ed.) Consequentialism and Its Critics (Oxford: Oxford University Press, 1988).
- Human Morality (Oxford: Oxford University Press, 1992).
- Schmidtz, David. “Islands in a Sea of Beneficence: Limits of the Duty to Rescue” Law and Philosophy 19. no. 6 (November 2000).
- Schneewind, J. B. Sidgwick’s Ethics and Victorian Moral Philosophy (Oxford: Clarendon Press: 1977).
- (ed.) Reason, Ethics, and Society: Themes from Kurt Baier (Chicago: Open Court, 1996).
- Sen, Amartya. Inequality Reexamined (Cambridge, Mass.: Harvard University Press, 1992).
- Resources, Value and Development (Cambridge, Mass.: Harvard University Press, 1997).
- Singer, Peter. “Famine, Affluence, and Morality” Philosophy and Public Affairs 1, no. 3 (Spring 1972): 229-243.
- Practical Ethics 2nd ed. (Cambridge: Cambridge University Press, 1993).
- “Outsiders: our obligations to those beyond our borders,” in Deen K. Chatterjee (2004).
- Slote, Michael. Common-Sense Morality and Consequentialism (London: Routledge & Kegan Paul, 1985).
- Smart, J. J. C. and Bernard Williams. Utilitarianism (Cambridge: Cambridge University Press, 1973).

- Stocker, Michael. "The Schizophrenia of Modern Moral Theory," The Journal of Philosophy 73 (1976): 453-66.
- Thompson, Michael. "What is it to Wrong Someone? A Puzzle about Justice," in Wallace, et. al. (2004).
- Toumani, Amadou and Blaise Compaore, "Your Farm Subsidies Are Strangling Us," Op-Ed, The New York Times July 11, 2003, A19.
- U.S. Committee for Refugees, World Refugee Survey 2001 (Washington, DC: U.S. Committee for Refugees, 2001).
- Unger, Peter. Living High and Letting Die: Our Illusion of Innocence (Oxford: Oxford University Press, 1995).
- United Nations, 'We the Children: End-decade review of the follow-up to the World Summit for Children,' Report of the Secretary-General, A/S-27/3, United Nations, New York, 4 May 2001.
- Walker, Margaret Urban. "Partial Consideration," Ethics 101 (July 1991): 758-74.
- Wallace, R. Jay. (ed.), Reason, Emotion, and Will (Aldershot: Ashgate, 1999).
- . Responsibility and the Reactive Sentiments(Cambridge, Mass.: Harvard University Press, 1994).
- . Philip Pettit, Samuel Scheffler, and Michael Smith (eds.). Reason and Value: Themes from the Moral Philosophy of Joseph Raz (Oxford: Clarendon Press, 2004).
- . "The Deontic Structure of Morality," forthcoming.
- Williams, Bernard. "A Critique of Utilitarianism," in J.J.C. Smart and Bernard Williams (1973): 77-150.
- . Ethics and the Limits of Philosophy (Cambridge, Mass: Harvard University Press, 1985).
- Wolf, Susan. "Moral Saints," Journal of Philosophy vol. 78, no. 8 (August 1982): 419-39.