DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY
Office for Scholarly Communication

Duties of Rescue: a Moderate Account

Citation

Nishimoto, Craig Takeshi. 2013. Duties of Rescue: a Moderate Account. Doctoral dissertation, Harvard University.

Permanent link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:11181182

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story .

Accessibility

DUTIES OF RESCUE: A MODERATE ACCOUNT

A DISSERTATION PRESENTED

BY

CRAIG TAKESHI NISHIMOTO

TO

THE DEPARTMENT OF PHILOSOPHY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHILOSOPHY

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS

September 2013

© 2013 Craig Nishimoto

All rights reserved.

Duties of Rescue: A Moderate Account

ABSTRACT

Peter Singer's famine-relief argument forces attention upon a class of requirements that make demands of an interesting kind. These are requirements to rescue another when hardly anything is at stake for the would-be rescuer and yet the utmost is at stake for the would-be rescuee. It is common to regard such requirements as moral duties, with all the significance that this implies. Since, however, the circumstances in which such duties apply can apparently repeat ad infinitum, the cumulative costs of compliance threaten to become shockingly great. I argue that there is a deep and legitimate challenge here that should neither be dismissed nor combined with superficially similar puzzles. Rather, for those of us wishing to maintain a comfortably moderate view about the demands of morality, the challenge is to find something to say that doesn't smack of mere rationalization. It is the kind of challenge that, when duly addressed, is likely to force broader changes in the way we think.

Broadly speaking, my method is to give successive clarifications of the challenge interspersed with attempts to locate and develop the most promising conservative response to it. Clarifications of the challenge inform the response, which in turn allows refinements to the challenge. The result is a clarified version of the famine-relief argument along with a comfortably moderate response that, while it asks us to accept

certain broader implications, recommends those implications as crucial and plausible alternatives to revisions that would, I expect, seem to many more radical. Rather than creating problems elsewhere, the recommended implications may provide useful perspective on a variety of moral phenomena, including the general relevance of compliance costs, the nature of deontological constraints, the category of special duties, and the justification of consequence-based evaluations of actions when the individual effects of such actions are negligible.

TABLE OF CONTENTS

Chapter One: The Problems of Demandingness…… 1

1. The Demandingness of Morality Generally…….. 1
2. The Demandingness of Moral Theories……. 5
3. Self-Standing Moral Demandingness……. 9

Chapter Two: The Limitations of the Famine-Relief Argument…… 15

1. The Famine-Relief Argument and Its Weaknesses……. 15
    - 1.1 The methodological challenge
    - 1.2 The grounding challenge
    - 1.3 The disanalogy challenge
    - 1.4 The empirical challenge
2. Efforts to Improve the Famine-Relief Argument…….. 30
    - 2.1 Unger's attempt
    - 2.2 Cullity's attempt

Chapter Three: The Differentiation Question…….……... 39

1. The Differentiation Question…… 39
2. Pressing the Differentiation Question…….. 43
    - 2.1 A general duty-specifying principle
    - 2.2 Appeals to analogy
    - 2.3 Appeals to analogy and skepticism of the will
    - 2.4 Side-stepping the empirical challenge
    - 2.5 Taking stock
3. Are We Committing Some Fallacy? …… 58
    - 3.1 Iterating when we should aggregate?
    - 3.2 A Sorites paradox?
    - 3.3 Like the paradox of the self-torturer?

Chapter Four: The Pond Question and the Famine Question…… 71

1. The Pond Question…….. 72

1.1 Sufficient answers and inappropriate additions
1.2 Avoiding Fishkin's inference

2. The Famine Question…….. 80

2.1 The centrality of the costs of repeated rescue
2.2 The student's progression of thought
2.3 A normative lens
2.4 The pond asymmetry
2.5 Applying the constraints
2.6 The alignment test
2.7 The possibility of more ideal collective responses
2.8 Feinberg on the relevance of the more ideal collective solution
2.9 Rawls's division of labor
2.10 Herman on inherited obligations
2.11 Respecting the priority
2.12 Respecting the content
2.13 Murphy on responsibility only for one's fair share
2.14 Conclusions

3. The Composite Question……108

3.1 The promising shape of a moderate solution
3.2 Recasting the challenge
3.3 Biting the first bullet
3.4 Biting the second bullet
3.5 Conclusions

Chapter Five: The Stringent Duty of Rescue……. 129

1. A Neglected Question…… 129

2. The Challenge of Stringency…….. 132

2.1 Priority
2.2 Priority and special duties
2.3 Significance
2.4 Direction

3. From the Millerian Value to a Stringent Duty…….…….. 145

3.1 Unacceptability judgments
3.2 To practical norms applying to individuals
3.3 Adding deontic direction
3.4 Explaining deontic priority

4. The Pond Question Revisited…….…… 165

5. The Differentiation Question Revisited……. 168

6. Rescue Duties as Special Duties……...……. 172

6.1 Responding to Scheffler
6.2 The category of special duties
6.3 Dismissing the over-demandingness loophole

7. Conclusions……...… 184

Bibliography……..... 186

ACKNOWLEDGMENTS

I thank Tim Scanlon and Frances Kamm for conversations and feedback on this project and on all of its drafts. I am especially grateful to Selim Berker for inquiring after me and joining my committee while I was out of residence. I have greatly profited from his careful eye and exceptional ability to readily engage my ideas, both written and spoken.

On separate occasions I presented content from chapters 2-5 to the Workshop on Moral and Political Philosophy at Harvard. Many of my arguments were constructed with an ever-optimistic eye towards convincing this famously critical audience. The participants provided helpful comments, and I am grateful to Nico Cornell, Johann Frick, Micha Glaeser, Kristi Olson, and Jeremy Fix for follow-up conversations.

I thank Mathias Risse for highlighting to me Richard Miller's contribution, and Derek Parfit for his encouragements about the importance of the topic. Christine Korsgaard turned my interests towards moral philosophy, and to duties of rescue in particular. Reaching further back, I again thank Andrew Hsu. I want to thank Jim Scheer, my selfless mentor, for his encouragements towards academic pursuits when my own mind was elsewhere.

I wrote most of this dissertation in Chapel Hill, where generous members of the UNC department welcomed me into their community. I am especially grateful to Adam Cureton, Doug MacLean, Susan Wolf, Tom Hill, Ryan Preston-Roedder, Geoffrey Sayre-McCord, and Kiran Bhardwaj. I thank Rob Rempher for always keeping his home open to me late.

I owe my deepest debt of gratitude to my family, whose unquestioning support for my work made it possible. I have lost count of the cross continental trips made by my

parents and by Eugene and Irena so that I might study with the full assurance that my children, now nearly four-years old, were superbly cared for. My debts to Katya, my eminently reasonable partner, are too numerous to list.

I dedicate this dissertation to Sarah VanWagenen, our cherished friend. Amidst several concurrent pursuits she has been writing her thesis, in linguistics and neuroscience, while I have worked on mine. If not for her struggle against cancer we would have finished these dissertations together. My argument defends the permissibility of a comfortably moderate position, but I know of no one who better demonstrates the beauty of casting aside moderation in exchange for richer ideals and fuller attentiveness to others.

For Sarah VanWagenen

whose aspirations lend honor to this pursuit

CHAPTER ONE:
THE PROBLEMS OF DEMANDINGNESS

1. THE DEMANDINGNESS OF MORALITY GENERALLY

Philosophers commonly think about the demandingness of morality, but not about the demandingness of self-interest, personal projects, aesthetic ideals, sportsmanship, philosophical rigor, etc. This focus plausibly arises from four properties that characterize moral requirements, each inferred from the way we ordinarily think about such requirements, and from the role of moral judgment in our practical deliberations and our assessments of actions and people. These four properties are priority, importance, categorical application, and pervasiveness.[1]

The *priority* of moral requirements is their status relative to other sorts of practical considerations. We commonly associate this feature of deliberative priority with the idea that moral requirements are overriding, or decisive. That is, if an action is morally required, then this is often assumed to give us decisive reason to perform the action—regardless of any and all our reasons for doing otherwise. If faced

[1] I draw the ideas of priority and importance from Scanlon 1998: ch. 4, categorical nature from Kant,

with the choice of either violating a moral requirement or sacrificing some cherished personal goal, we suppose that the only correct decision, ordinarily, is to sacrifice the personal goal. Even if weighty reasons of personal interest, aesthetics, sportsmanship, special relationships, etc., all combine against fulfilling a moral requirement, the latter requirement is commonly thought to carry the day, overriding or undercutting every consideration to the contrary.

The *importance* of moral requirements concerns the special significance appropriately attached to compliance with moral requirements. Though we can overlook some deviation, we pay special attention when someone neglects or disregards moral considerations, especially when this is done deliberately, broadly, or somewhat gratuitously. The comparable neglect or disregard of nonmoral considerations does not, at least in general, occasion the same sort of serious concern and censure. Comparatively, violations of moral requirements tend, more generally, to license attitudes of guilt, remorse, resentment, indignation, and blame. When a violation of a moral requirement takes place, this is moral typically something that has to be addressed.

The *categorical application* of moral requirements concerns the way in which a moral requirement upon a person does not depend on that person's contingent aims. While this property isn't unique to moral requirements (Foot 1972), it does mark a way in which moral requirements are less escapable than others, a way in which a person has less control over whether or not a moral requirement applies to him or her. So, while the requirements of tightrope walking may be severely confining, they are more obviously self-imposed and optional since its requirements only apply insofar as a person voluntarily endeavors to attempt the feat.

Finally, the *pervasiveness* of moral requirement concerns the scope of moral evaluation: no type of voluntary action is in principle immune to becoming morally obligatory or forbidden. Even brushing one's teeth is susceptible to moral evaluation—and the activity may, for instance, be impermissible if one chooses to brush one's teeth while an infant drowns in the bathtub (Scheffler 1986: 533-34).

These apparent properties of moral requirements mean that such requirements are likely to conflict with other sorts of practical considerations. The pervasiveness and the categorical applicability of moral requirements mean that moral considerations can in principle invade any area of our lives. Combining these properties with priority and importance means that we are uniquely vulnerable to monopolization by morality, its demands confining us to a very narrow set of options, deviations from which are failures of an especially serious kind.

In addressing concerns about demandingness, it may be appropriate to examine whether or not we really should ascribe these properties to morality.[2] My immediate aim, however, is to think about how such apparent properties can give rise to separable

[2] Since I will be investigating Peter Singer's famine-relief argument, it is worth nothing that Singer doesn't assume the priority of moral requirements (1999a: 308, 2007: 478-79, 82). While this may remove certain grounds for opposing Singer's conclusion in his famine-relief argument, it does so at the cost of his conclusion's significance (Singer suggests that reasons arising from "[one's] own interests, and those of [one's] family" may regularly "counteract the demands of morality"). If the priority of moral requirements is implied by ordinary usage of the relevant moral judgments, then Singer would be making room for the revisionist demands of his particular conclusions through a revisionist relaxation of the idea of moral obligation generally. If I am right about ordinary usage (that moral obligation has the property I am calling "priority"), and if Singer isn't committed to simply changing ordinary usage (and why should he be?), then he can state his intended conclusion as follows. "While the requirements of 'Singer-morality' are highly difficult to fulfill with regard to famine victims, we shouldn't assume that such requirements are, in the usual and most general sense, obligatory." Singer, however, would likely want to insist further that the requirement to act on behalf of distant famine victims has (relative to considerations arising from one's own interests and those of one's family, etc.) *as much priority* as the requirement to rescue the nearby child from the shallow pond. Granting that the requirement to rescue the nearby child garners the high priority usually attributed to moral obligation generally, Singer's rejection of the priority of moral obligation generally is a moot point, despite his suggestion to the contrary.

concerns about demandingness. This will help me to distinguish the particular problem of demandingness I intend to address.

By itself, the supposed priority of moral requirements can prompt the question why we should grant it that status. Combined with categorical applicability, moral requirements are typically thought to constitute decisive reasons even for the person who sees no point in acting in any of the particular ways that morality prescribes. The task is to make better sense of this authority. If, warned by H.A. Prichard against appealing to non-moral values, all we can say is that this way of acting is right, or that it is one's duty, even normal people might sympathize with the amoralist's failure to be moved (1912; Scanlon 1998: 149-50). Such answers leave the demands of morality mysterious.

So this first kind of "problem" of demandingness might be described as the problem of explaining morality's authority. I take this problem primarily to be one of finding a satisfying explanation; it is the task of providing an account of morality that renders more intelligible the legitimacy of these commonly ascribed features. A satisfying explanation here is likely also to contain an explanation of morality's importance. That is, a good explanation is likely to explain the priority of morality in terms of a grounding, or a subject matter, that ordinary people can readily appreciate as significant in those distinctive ways that are characteristic of morality's importance.

A philosophical account that addresses this first broad problem of morality's demandingness isn't likely to be neutral between competing moral frameworks and first-order moral judgments (Scanlon 2003a: 129). If, for example, the Kantian appeals to the rational requirements of agency to explain the priority and categorical applicability of a moral requirement, this tends to lead to distinctively non-consequentialist principles, such

as a prohibition against the lying promise derived from the formula of universal law through the "contradiction in conception" test (Kant, Gregor 1998: Section 2). On the other hand, when the consequentialist seeks to explain the general priority of morality in terms of the well-being of people, or the idea that "it can never be right to prefer a worse state of affairs to a better" (Foot 1985), then teleological principles are likely to displace the kinds of constraints characteristic of Kantian views.

Insofar as the moral framework points us to the grounds from which we might hope to account for morality's distinctive authority, those same grounds are apt to serve as the guidelines for identifying more particular moral requirements. Whichever theory we choose, it is likely to give rise to a second kind of demandingness problem, wherein we must deal with any apparently excessive requirements favored by that theory. If, for example, we adopt as a foundational thesis the idea that it can never be right to prefer a worse state of affairs to a better one, then this will plausibly favor a prohibition against purchasing new shoes, whenever something even slightly better can be done instead.

2. THE DEMANDINGNESS OF MORAL THEORIES

This second kind of demandingness problem, therefore, concerns views about the demands of morality reached through the conscious application of a theory. Sometimes the theory is quite minimal, as when Thomas Nagel proposes that moral judgment requires reconciliation with an impersonal point of view (1995, 1986). Demandingness problems arising from more highly specified theories are easy to find.[3]

[3] The demandingness problems associated with consequentialism are the most well-known (Williams 1973, Kagan 1989, Scheffler 1994). Williams (1981) discusses demandingness objections to Kantianism;

While it is the conscious application of the theory that generates the problematic demands in this kind of demandingness problem, it isn't always clear what generates the resistance to those demands. Often the resistance to the demands seems to arise independently of any conscious application of principles. In that sense, then, the resistance to the demands is *intuitive*.[4] This isn't to say that one always arrives at an intuitive response without careful consideration of the subject matter; it is only to say that a particular act of consideration isn't necessary, namely the conscious inference from general principles to that intuitive response, or the conscious recollection of having made some such inference. If we are asked to defend our intuitive responses, we can often give reasons, but in formulating these justifying reasons, the task will typically be to newly locate or formulate them, rather than to remember them. The intuitive response may be a belief or judgment, or it may be some reactive attitude. The objects of such responses are also diverse. They may be directed towards specific attitudes or actions or other features in a specific circumstance. If the intuitive response is a belief, it may be a belief about the moral status of a given action or attitude, the moral significance of a given factor, the sufficiency of a given reason, or the legitimacy of a general principle.

To say that one's resistance to the demands of given moral theory is intuitive isn't to say much. A judgment that I now arrive at intuitively might once have been a conclusion I would reach only after a consciously processed inference from principles that I later "internalized." Intuitive judgments or inclinations are also commonly false or

---

Ashford (2003) argues that Scanlon's contractualism is as demanding as utilitarianism when it comes to helping others meet their basic needs.

[4] I borrow the phrase "conscious application of principles" from Frances Kamm's characterization of intuitive judgments. Kamm similarly contrasts intuitive judgments with those that are "deliberately constructed" (2001: 11-12).

misleading. It is sometimes easy to debunk them, to explain why we have them in a way that provides no promise for justifying them. If the intuitively arrived at belief is inconsequential with respect to one's other beliefs and commitments, it may be easy to repudiate. As a child, I had persistent, idiosyncratic, and conflicting convictions about compass orientation. I believed that the large window of my first-grade classroom faced west. Google Maps now shows that this belief was false, but it might have been debunked early on by noting the orientation of my classroom desk from which I would have viewed all the maps.

The cases that are of more general interest will involve intuitive moral responses that we can reasonably expect others to share, or at least sympathize with. They will involve that subset of intuitive responses that belong to what is called commonsense morality. As the term is currently used, "commonsense" designates a privileged class of commonly shared, or at least peculiarly shareable, beliefs.[5] Mere popularity isn't enough, for even if some religious dogma, such as the Immaculate Conception, gained overwhelming popularity through demographic change, perhaps we should not say that it becomes a commonsense belief. What accounts for the privilege is not the popularity of commonsense beliefs, but rather their special kind of plausibility.

This kind of plausibility we attribute to commonsense beliefs explains why they tend to be commonly shared, or are at least peculiarly shareable. When we react with suspicion or incredulity at someone's denial of a commonsense belief, this is partly because the plausibility of such beliefs is thought to be independent of one's acceptance

[5] Though I here speak of commonsense *beliefs,* a broader range of moral attitudes, views about morality, and perhaps forms of moral reasoning could easily be said to belong to commonsense morality. We might say that, in a given circumstance, responding with a certain reactive attitude is part of commonsense morality (along with the judgment that such reactions are appropriate).

or non-acceptance of any special theory or ideology. We expect, therefore, that any person of normal capacity can be brought around to at least appreciating the intuitive plausibility of a commonsense belief, even if this appreciation requires a special effort of attentiveness on their part, along with care to avoid certain of the defects mentioned above.

If we are dealing with what we regard to be a commonsense belief, then someone's failure, or refusal, to accept that belief tends make one suspicious of defect in that person. We suspect a misunderstanding, insufficient attention to (or ignorance of) key facts, false beliefs, perceptual malfunction, imbecility, insanity, or some kind of mental mismanagement. Or, in those cases in which it is unreasonable to suspect defect in the person in question, such a one is still subject to our incredulity, and this sentiment often lingers even if that person's rejection of the commonsense belief is seen to make sense in light of his or her otherwise respectable belief system.[6]

Commonsense beliefs tend to be the kinds of beliefs that are simply taken for granted. They tend to be the kinds of beliefs by which we evaluate other ideas. Hence, the second demandingness problem is often used as a criticism of the theory at issue. Often defenders of the theory will try to show that there isn't any conflict with commonsense morality after all, or they will try to modify the theory so as to make this so. When one instead opts with the theory against some aspect of commonsense

[6] Some examples might help. To commonsense morality belongs the belief that one should not gratuitously inflict suffering upon random strangers, and also the belief that, if it costs you nothing to prevent a horrific traffic accident, you probably should. These principles strike me as particularly plausible and I would expect others to share them (barring deep disagreement over relevant facts, such as the fact that preventing horrific traffic accidents would not generally frustrate the workings of a perfectly good and omniscient god). If I were to encounter someone who did not at least appreciate the plausibility of such principles, then I would suspect some deficiency in that person or in my communication of these principles. If these and other defects were ruled out, and the person still couldn't in the least even sympathize with my sense of the plausibility of these principles, I would be astonished.

morality, the revision tends to be profound, for what gets revised are beliefs that are broadly accepted and particularly plausible. They are beliefs around which we have shaped other ideas and commitments. These last points suggests why challenges to commonsense morality are of particular interest whether or not the tenets of commonsense morality really should be accorded their de facto privileged status.

3. Self-standing moral demandingness

Much of what gets called a demandingness problem fits this characterization of the second demandingness problem, where the problematic demands depend on principles distinctive of some particular theory about morality. Are there, however, problematic apparent moral demands that arise independently of particular moral theories, or which are at least common to a broad range of moral theories? If an intuitively apparent moral requirement is problematic whenever it is in tension with other pressing practical considerations, then there would be many problems of this sort. Many hard cases, apparent moral dilemmas, and puzzles involving trolley cars would fit this characterization. However, the problems I am most inclined to think of as problems of demandingness have a further feature: resistance to the demand appears to arise, at least in part, from considerations of self-interest or partiality, as opposed, in particular, to competing moral considerations. So, in a paradigmatic demandingness problem, the tension is between some apparent moral requirement and some other consideration of special importance to the person to whom the apparent moral requirement would apply. That is, when I face one of these demandingness problems, the apparent moral

requirements pull against my own self-interests, or the special projects or relationships that happen to matter to me. Moreover, while I can imagine many scenarios in which the demands of morality would conflict with important considerations of these kinds, I think of demandingness problems as those cases that call into question the legitimacy of the apparent moral demand. If a problem can be found that fits these constraints, then there is a third kind of demandingness problem. Unlike the first, it is a problem internal to morality. Unlike the second, the problematic demands are largely independent of particular moral theories.

A prominent candidate for a demandingness problem of this third sort is the kind of challenge pressed by Peter Singer in "Famine, Affluence, and Morality" (1972). Singer's intent is to present an argument that "does not depend on any specific values or ethical principles" (2011: 202), and, if successful, his argument would show that it is wrong to do many things that most of us regard as entirely permissible: spending on non-essential consumables, and pursuing most hobbies or any kind of comfortable, non-lucrative career.

If something like Singer's argument does provide a demandingness challenge of the third kind then it is worth treating it as such, rather than regarding it as a problem raised by utilitarianism or of some particular moral theory. When, as is often done, the problem is subsumed under the second demandingness issue, fundamental tensions in our widely shared assumptions may get buried under problems of a more manufactured sort. By embedding the demandingness problem in a particular theory, we are limited by the resources of that theory when formulating the problem, even if such resources only poorly capture the relevant tensions. Free from such restrictions, we may discover that

the pull towards highly demanding conclusions is not most convincingly represented as requirements to prefer the better state of affairs, to promote individual well-being, or even any general obligation of beneficence. Reliance on a theory may also constrain our resources for resolving the problem. The rescue cases may seem to be just one more case for which agent-centered prerogatives must be explained, or for which the personal point of view needs to be accommodated. Such efforts may only take us a step backwards to work on a problem whose resolution doesn't promise to resolve the specific tensions of the rescue case. We may also find that any overdemandingness objection framed on behalf of our intuitive judgments fails to get off the ground. Proponents of utilitarianism, for example, can respond that objections about over-demandingness simply beg the question against the theory (Kagan 1984, 1989, Murphy 2000, Sobel 2007).

Are there other demandingness problems of the third sort? I am not aware of any. More exactly, I don't know of any such problems that are nearly as compelling as those that, in the spirit of Singer's famine-relief argument, appeal to individual duties to rescue particular others from serious peril when this can be done at little cost to ourselves.[7] If

[7] Garrett Cullity suggests that there are other demandingness problems of the third kind. He proposes that all plausible moral views face challenges from the famine-relief argument because they will all "make the promotion of others' welfare at least a *part* of morality" (2009: 9). Similarly, Cullity suggests, all plausible moral views will recognize other aspects of morality that give rise to additional demandingness problems. He proposes two such aspects: (a) that "minimal restrictions on my liberties can be required in order to protect others' liberties against great curtailments," and (b) that I can be required to contribute to cooperative activities that protect or produce public goods. Both aspects, according to Cullity, can, through repeated application, result in an accumulation of small constraints or requirements that, in aggregate, constitute demands far greater that many would find tolerable. However, for (a) to generate intolerably great demands the world may have to change a great deal (I give an example of such a change in chapter four). The curtailments Cullity does suggest strike me as either reasonable and non-threatening, or they fail to be examples of one person's behavior greatly curtailing another person's liberty (16). As for (b), Cullity himself recognizes that the needed versions of the relevant principle are controversial (18). I can plausibly deny that I am required to contribute to every cooperative activity that supports important public goods that I enjoy. In a proliferation of such activities and public goods, perhaps I contribute my fair share to the general public good when I pick some cooperative activities to support and I try to avoid obstructing the others.

such problems are unique, then I suspect it is because they exploit what can be seen as a unique vulnerability in a standard taxonomy of moral duties. Consider the following divisions of moral duties.

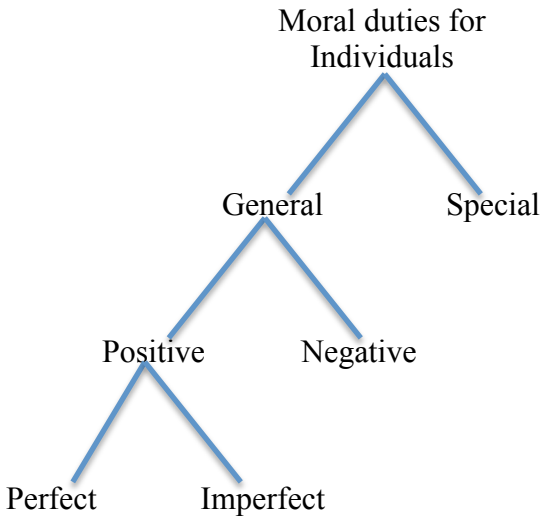

The categories represented here are familiar.[8] What is noteworthy is how each of the rightward branches represents a way of limiting the demandingness of moral duties falling to individuals. Special duties, as opposed to general ones, are only owed by those who have made a special, voluntary transaction, or to those with whom one stands in some special relationship. Duties of the latter kind are limited because one naturally only stands in the relevant relationship with few people (one has only so many parents or siblings). Duties of the former kind are only those that have been, in some sense, voluntarily assumed. One can regulate how many, or what kinds of, promises one makes, or which roles one adopts; one doesn't have to enter into a friendship, a marriage, or a

[8] Rawls presents a take on this taxonomy in sections 18-19 of *A Theory of Justice*, without including the division between perfect and imperfect duties. H.L.A. Hart's distinction between special and general moral rights is often generalized to distinguish special and general obligations (1955: 183). The distinction between perfect and imperfect duties has been variously characterized and traces back to Grotius and Pufendorf, and through Kant and Mill (Schroeder, forthcoming).

special contract with everyone, and when one voluntarily enters into such roles or make such contracts, the resultant demands are commonly offset by various benefits (the demands of a friendship, for example, go both ways). A negative duty ("Never kill") tends to be less monopolizing than its positive correlates ("Always prevent killings"). Imperfect duties characteristically allow the agent significant latitude in how they are to be fulfilled. A prototypical duty of this sort issues a demand like this: help someone or another, at sometime or another, in some or another way. Imperfect duties are therefore the kind of duty you can fit into a schedule.

If we follow the leftward branches all the way down, we come to a subset of moral duties for which there is no category-given means of limiting demandingness. These are the general, positive, perfect duties. Within this category, duties will vary according to the significance we would attach to their violations. So, while I might have a duty to helpfully give directions to a passing driver from out-of-town, violating such a duty would not be as significant as refusing to rescue a drowning child from a shallow pond. If the famine-relief argument is uniquely compelling among demandingness problems of the third sort, it is because of the way it leverages a particularly serious moral duty within a category of duties that uniquely lacks a means for limiting demands.

The famine-relief argument may well be a uniquely compelling example of a demandingness problem of the third type. Moreover, if the pull towards its problematic demands is securely grounded in judgments of commonsense alone, then, for the reasons suggested earlier, our rejection of those demands will likely require profound revisions to widely shared assumptions.

If such revisions are to be made, a convincing case for the problem will have to precede them. Singer's own presentations of the famine-relief argument can be criticized on several grounds. I do not believe that his arguments secure his highly demanding conclusions independently of controversial moral and empirical premises that go well beyond the kinds of intuitive moral judgments that I am interested in here. In the next chapter, I survey the limitations of the famine-relief argument along with attempts to improve it.

Chapter two's clarifications set the stage for the construction of a challenge, in chapter three, that draws upon the main elements of the famine-relief argument while working within its limitations. There I provide a demandingness problem of the third kind where I try to show the ways in which the pressure towards the problematic demand is rooted in commonsense morality.

As I examine the famine-relief argument and seek to provide my own version of it, my ambition is always to better understand what it is that makes the argument compelling, and what it is that makes it so difficult to provide a satisfying rejection of its demanding implications. It is in the clarification of the challenges that I look for hints concerning their most promising resolution. The hints I uncover lead me, in chapters four and five, to what I take to be the most promising defense of a moderate position of our moral duties. The moderate position will itself simply be the rejection of the famine-relief argument's shocking implications, and I will be looking for defense of that position that is both philosophically satisfying and minimally revisionist. When I must make revisions, I try to show them to be the necessary costs of opposing revisionistically demanding views about duties on behalf of strangers.

CHAPTER TWO:
LIMITATIONS OF THE FAMINE-RELIEF ARGUMENT

In this chapter I consider Peter Singer's famine-relief argument and its limitations. I then consider the limitations of two of the most developed attempts to improve Singer's argument. All this sets the stage for clarifying an important challenge to which Peter Singer and others have drawn attention.

1. THE FAMINE-RELIEF ARGUMENT AND ITS WEAKNESSES

Arguments for the claim that we as individuals are obligated to make shockingly great sacrifices on behalf of strangers typically appeal to some *general duty-specifying principle* that specifies something individuals are morally required to do. Singer proposes, "[I]f it is in your power to prevent something bad from happening, without sacrificing anything nearly as important, it is wrong not to do so" (2009: 10). Though Singer's principle specifies duties to prevent bad events, alternatives more narrowly require individuals to help others, or to rescue others from dire straits. To the general

duty-specifying principle is added further claims about how we continually face scenarios for which the principle applies. The conclusion is then drawn that our moral duties are shockingly demanding. Following Singer, I will call any argument of this form a *famine-relief argument*.[9]

While we can find problems with each version of the argument, I find certain of Singer's general strategies promising. I will formulate several challenges that draw from Singer's own arguments and the attempts of others to improve upon them. My goal is not the ambitious one of trying to secure Singer's demanding conclusion; it is rather to show only how difficult challenges arise from our own intuitive or commonsense moral judgments.

Let us first survey the main weaknesses in Singer's attempts to secure a highly demanding conclusion. I count four. Three are problems with Singer's strategies for securing his general duty-specifying principle. The fourth weakness concerns the empirical claims required to secure the demanding conclusion from such a principle.

1.1 The methodological challenge

[9] Singer has defended the famine-relief argument (so named in Singer 1977), with minor variations, since "Famine, Affluence, and Morality" (1972). Here is a recent formulation:

> First premise: Suffering and death from lack of food, shelter, and medical care are bad.
>
> Second premise: If it is in your power to prevent something bad from happening, without sacrificing anything nearly as important, it is wrong not to do so.
>
> Third premise: By donating to aid agencies, you can prevent suffering and death from lack of food, shelter, and medical care, without sacrificing anything nearly as important.
>
> Conclusion: Therefore, if you do not donate aid to aid agencies, you are doing something morally wrong. (2009: 15-16; see also 2011: 200-12, 1972: 231)

Singer's first strategy is to select a general principle of moral duty that is intuitively plausible.[10] The basic problem here is there is ample room for skepticism when the object of our intuitive judgments is a principle, and especially when the principle is both highly general and is designed to carry shocking implications. When a principle is intuitively judged to be correct, this may only be because the principle is imprecise and malleable (the "Golden Rule" and Kant's Humanity Formula are, I think, good examples), or because one hasn't yet fully considered the principle, failing to realize certain of its problematic implications (as with Frege's Basic Law V). Although it is dangerous to conclude that a principle is true before appreciating what the principle commits us to, this often seems to be exactly what we do when forming an intuitive judgment about substantive principles (McMahan 2000). In various proposed versions of the famine-relief argument we can readily see a failure to duly assess the implications of the proposed general principles of individual moral duty. When these principles are formulated so as to generate Singer's intended conclusion, their plausibility is often superficial. Scrutiny tends to show that the proposed principle is more dubious than initial appearances suggest, requiring tweaking or fine-tuning that render the principles either too weak or inexact to sustain the original conclusion, or too complex to sustain their intuitive appeal.[11]

[10] Singer is sometimes interpreted as inferring the general principle from cases in which one is intuitively required to rescue someone, e.g., Gomberg 2002: 31. Singer does not interpret his argument as relying on this dubious inference.

[11] For variations of the relevant principle see Sidgwick 1981: 348-49, Singer 1972: 231, 235, Slote 1977: 127, Arthur 1977: 48, McKinsey 1981: 313, Fishkin 1982: 65, Unger 1996: 57-58, 144, Scanlon 1998: 224, Herman 2001: 227, Kekes 2002: 514, Ashford 2003: 287, Cullity 2004: 19. For principles that are clearly too weak (often by design) to easily secure Singer's conclusion see Unger 1996: 12, McGinn 1996, Hooker 2002: 167, Miller 2004: 359, Cullity 2004: 174.

Insofar as the intuitive attractiveness of a principle is vulnerable to counterintuitive implications attributable to the principle, Singer's style of argument threatens to undermine itself: if the argument is successful, then it shows that the general principle of beneficence has surprising and counterintuitive implications; these same implications, however, simultaneously threaten the initial plausibility of the general principle of beneficence. While threats of this kind can be overcome in various ways, it typically requires work to do this (to somehow show, for example, that the problematic implications of the principle are more acceptable than the problematic implications of rejecting it). Garrett Cullity has called it the "methodological challenge": if the key principle's intuitive plausibility is all that it has going for it, then its strongly counterintuitive conclusions substantially threatened it (1994: 104-5, 1996: 52). By simply giving greater credence to our intuitive resistance to Singer's demanding conclusion, some would view Singer's argument as a reductio ad absurdum of his principle (Trammell 1975: 133-4, Lewis 2000: 155, Thomas 2009: 132).

1.2 The grounding challenge

Singer's second strategy is to reinforce the plausibility of his general principle of beneficence by grounding it in a more fundamental feature of morality. Singer tries to ground his principles in a fundamental requirement, named "impartiality" or "universality," which he claims has been recognized throughout history, across cultures, and by a wide range of moral theorists (1977, 2004, 2011: 10ff). In the widely varied traditions that Singer cites, it is far from clear that there is convergence on any universalization requirement suitable for his purposes (in the Judeo-Christian

formulations of the "Golden Rule," in Kant's formula of universal law, in the impartial spectator of Hutcheson, Hume, and Smith, in the basic axioms of the utilitarians, and in Rawls's use of the veil of ignorance). If there is convergence in these traditions, it is it hard to believe that it includes the universalization requirement that Singer claims is most precisely formulated by R. M. Hare (1995: 174, citing Hare 1965: ch. 2). Singer's own formulations of that requirement are dubious.[12]

Consider, for example, Singer's formulation in terms of giving "equal consideration" to interests, regardless of to whom they happen to belong.

> The essence of the principle of equal consideration of interests is that we give equal weight in our moral deliberations to the like interests of all those affected by our actions. This means that if only X and Y would be affected by a possible act [and its alternatives], and if X stands to lose more than Y stands to gain, it is better not to do the act. (2011: 20)

There is a sense in which it is better to act as such a principle advises. This would the sense in which more interests, or more weighty interests, would be met. It is difficult, however, to find any uncontroversial way to translate this into a standard for moral permissibility. Even if we accept that impartiality is an important feature of morality, we needn't accept the idea that, at a fundamental level, the moral point of view bars the individual from assessing options with partiality.[13] So, although such a principle might ground Singer's preference-utilitarianism, it is at best a highly controversial interpretation

[12] Singer nearly concedes these points in *Practical Ethics* (2011). Instead of deriving his utilitarianism from the widely recognized "universal aspect of ethics" Singer instead proposes that it establishes a presumptive case for "a broadly utilitarian position," such that good reasons need to be given for any departure from utilitarianism (11-12). Insofar as the famine-relief argument is designed to appeal to those who believe there are good reasons for departing from utilitarianism, Singer's maneuver here is ineffective.

[13] Singer notes: "Kamm correctly states that, in contrast to the views held by Samuel Scheffler and Thomas Nagel, my conception of the moral point of view does not, at the critical level, allow us to weigh things partially" (Singer, Jamieson 1999: 308).

of what morality is about. It is too controversial to add much plausibility to the famine-relief argument's key principle.

Besides this equal consideration standard, Singer also proposes a striking procedure in his discussions about universalizability:

> As I have argued elsewhere, I know of no better way of deciding what we ought to do than to put ourselves in the position of all those affected by the actions open to us and do what we would prefer if we had to live all their lives. (2010: 195)

One might interpret this as mere advice for how we might improve, in a moral respect, our decision-making.[14] Such an imaginative exercise could be helpful in the way that taking a deep breath and counting silently to ten might be helpful before responding to someone in the heat of anger. Such tricks help one to overcome or counteract certain predictable problems in deliberation. To imaginatively put oneself in another's position is, I take it, a common way of raising one's awareness of morally relevant considerations, helping to counteract one's bias towards one's own interests. While I would have no objection to such advice, it fails to reveal the kind of fundamental feature of morality that Singer wants, the kind that could justify the principle of individual moral duty he needs for his famine-relief argument. Even if I can duly understand and appreciate the standpoints and conflicting concerns of every relevant person, any particular method of arbitration remains controversial. So while it may be morally helpful to consider the competing concerns of others, the helpfulness here presupposes, rather than provides, the moral standard for arbitrating the concerns of different people.

If Singer intends this suggestion to provide the content of "impartiality" or "universality" then we must reject it as unclear, incoherent, or unfeasible. Singer would be suggesting something akin to Rawls's original position, conditioned not by a veil of

[14] This reading may be suggested in Singer 1995: 174-75.

ignorance, nor by the prospect of having your position randomly or maliciously assigned, but rather by the certainty that we would actually live the lives of all of those affected "by the actions open to us." The problem is figuring out what it is that we are supposed to be imagining when we imagine living the lives of all these other people.

One thing that seems quite possible for me to do is to abstract away some of the particularities about myself and my own circumstance, and imagine that I instead share certain of the features and circumstances of another person. After doing this, I might further imagine what I, transformed in this way, would want or prefer. Not just any way of doing this will be adequate for determining moral guidelines. In selecting the features of myself to remove and the features of the other person to adopt, I have to make sure that I am adopting for the other person all the morally important features and shedding from myself all the features that would distort moral judgment, at least for the question at hand. This means we need a way of identifying the relevant features.

Since Singer's proposal doesn't provide guidance for identifying such features, it is tempting to read it in a more radical way: in placing myself in the other person's position, I am supposed to imagine losing *all* of the features peculiar to myself and my own circumstances and gaining *all* of the features of that other person and her circumstances. If, however, I try to imagine having only that other person's memories, attitudes, dispositions, relationships, plans, etc., then, to the extent that I am successful, I strongly suspect that this fails to be a case of putting *myself* in the other person's position.

When Singer suggests that we should put ourselves in another person's position and imagine having to live his or her life, it is not clear why we should expect the hypothetical preferences to helpfully correspond to anything in good moral deliberation.

If, on the one hand, I am completely successful in imagining myself with only other person's features, then my hypothetical preferences will simply be the preferences of that other person. All I am then doing is accurately imagining the other person from her own point of view. After successfully accomplishing this task, my resources for moral reasoning are no better than those of that other person who has yet to follow Singer's advice. If, however, I am less successful, then my hypothetical preferences will just be the preferences of a hypothetical person constructed from some mixture of my own features and the other person's features. It is difficult to see how it could be helpful to consider the preferences of a person who, for example, has my beliefs, memories, and friendships, but has some other person's tastes, intentions, and life goals.[15]

Suppose, however, that there is a way to work out such difficulties and that, in some relevant way, it is possible to put oneself in another's position. To follow Singer's advice for deliberation, we also have to figure out a sensible way of imagining how we are to do this for all of the relevant people. Since all the relevant people are all of those affected by the actions open to me, it is a great many people. We might initially suppose that we can limit the number, realizing that the kinds of actions open to me become limited by whatever actions I take. If I spend the rest of my life in Paris I cannot also spend the rest of my life in Malawi. If, however, my decision to spend the rest of my life in Paris implies that few Malawians are affected by the actions open to me, then Singer's formulation would bizarrely generate strong reasons against going to Malawi, or any

[15] In may be that Singer means that, in putting oneself in the position of others, one must only adopt the preferences of those others, as suggested in Singer, Jamieson 1999: 298. It is unclear, however, what this would mean if we had to put ourselves in the position of more than one person simultaneously. If Singer means that we give weight to others' preferences "in proportion to the strength with which they hold them" (1988: 152), or that we should "adopt the course of action most likely to maximize the preferences of those affected" (2011: 12), then this universalization requirement may boil down to Singer's principle of equal consideration, without removing any of our reasons for doubting that principle.

other place of relative misery. That is, if I had to live the lives of everyone affected by my actions, prudence would instruct me to guide actions according to *whom* I happen to affect: I should avoid affecting (in any way whatsoever) people with miserable lives; on the other hand, I should actively stalk people whose lives I would like to live, with the aim of affecting them in some trivial or positive way.

We must therefore interpret Singer's formulation in the following way: even if I decide to live in Paris, I must still imagine myself in the position of all of those whom I might affect were I instead to have decided to live in Malawi—or China, or India, etc. If, moreover, in any of these possible paths it would have been open to me to affect *in any way* those multitudes, then I must also imagine living the lives of all of those multitudes. Since for every person there is likely to be some course of action or other open to me through which I could affect that person (either intentionally or otherwise), then on Singer's proposal I must imagine living the lives of nearly all presently living people (I might perhaps exclude individuals who are just about to die in highly isolated places). Since, moreover, even minor effects on people who are alive now can alter who comes into existence later, I may also have to imagine living the lives of all people who are yet to be born.

While the number and diversity of all such people burdens the imagination, there is a more basic problem of figuring out how I am to imagine living multiple lives (even assuming, again, that there is an appropriately helpful way of imagining living one other person's life). One way of imagining living multiple lives, suggested by C.I. Lewis, is to imagine living each of these lives seriatim (1946: 546-47). We might try to imagine this as a long series of reincarnations back into the present world to live out one of the lives of

someone else. This innovation might seem to resolve certain problems associated with aggregation (small costs and benefits would become more significant to you as more people experience them because each new person represents a new cost/benefit to yourself). It may even do this without allowing trivial benefits to the many to outweigh extraordinary costs to the few (given that you would still have to suffer those extraordinary costs). Distortions, however, arise under the seriatum interpretation because of the way it suggests that I regard the lives of others as temporal segments of my own very long life. The problem with this is that it often makes sense for me to impose sacrifices during one time of my life for benefits that I will receive later. This means that the seriatum interpretation will tend to license the harming of one person for benefits received by others *as if* it is one and the same person experiencing the harms and receiving the benefits. Since, however, what is fair to a person so compensated for her sacrifices may not be fair to a person receiving no such compensation, the seriatum interpretation will distort our judgment.

A second interpretation, proposed as "highly speculative" by Thomas Nagel, is to imagine living all the lives "as a set of distinct individual lives….not in a series, but simultaneously" (1978: 140-42). This interpretation is supposed to help block the "interpersonal compensation for sacrifices" that the seriatum interpretation threatens to license. The cost, however, is apparent incoherence. To consider what should presumably be a relatively easy case, suppose that I only have to imagine simultaneously being both myself and one other person whom I know quite well. What is it to imagine myself simultaneously being both myself *and* my wife? The suggested task seems to be that of simultaneously imagining myself being just like myself *and* being just like

someone who is not like myself, and this seems a lot like trying to imagine a square circle. To imagine being just like someone who is not just like myself is to not imagine being just like myself.

Nagel argues, however, that

> the conditions of choice can be understood metaphysically, for we can imagine a person splitting into several persons, each of whom bears to the original, over time, the kind of relation that constitutes trans-temporal personal identity for an individual person….This provides a sense in which an individual might expect to become each of a number of different persons—not in a series but simultaneously—so that each of their lives would in a sense be his unique life…. (1978: 141)

If it is in this way possible to imagine becoming multiple persons existing simultaneously, then we might try to interpret Singer's proposal along similar lines. I am to imagine what I would prefer if I knew I were about to split into multiple persons, where these multiple persons are all those affected by the actions open to me. The value of Nagel's proposal is in showing that there is no conceptual barrier to imagining one's preference under the knowledge that one is about to become, in a sense, multiple and simultaneously existing persons.

The concern, however, is whether this element of fiction introduces any new distortions to moral deliberation concerning actions that affect people who aren't really future branches of myself.[16] Suppose I know that I will soon split into two persons in the way Nagel proposes, so that my present decisions will affect both branches. In making these present decisions, does mere prudence lead me to balance the interests of these two future branches in what would be a morally appropriate way were they not each branches of my present self? Under some circumstances, I think that the answer is "no."

[16] In thinking about the branching of a person, and of myself, I am indebted to Derek Parfit's discussions (1989: part 3).

I can imagine, for example, a scenario in which it would be intuitively prudent to take action that would severely sacrifice the interests of one of my future branches for the sake of another of my future branches. Suppose I am in need of a skin graft and I can issue orders to have myself split, or "duplicated," as I sleep, with the provision that either my duplicated self or the original will be immediately and painlessly killed for the sake of its skin. When only the one branch of me awakes the next morning, it will be just as if I have never branched, but I will have the needed skin for my graft. This strikes me as a potentially prudent alternative to removing healthy skin somewhere else on my non-duplicated body. I take it, however, that it would be wrong to kill any other person in his or her sleep for the sake of my skin graft.[17]

It may seem that a crucial factor is that the person to be sacrificed is a duplicate, or a near duplicate, of either the person whom the sacrifice will benefit, or of the person who has placed the orders to perform the sacrifice. Though this strikes me as plausible, it suggests a new problem with interpreting Singer through Nagel's proposal. If, instead of being able to nearly perfectly duplicate myself, I can instead only issue orders to create so poor a duplicate that it results in a person who bears no closer resemblance to me than any other random stranger, then it seems to me that there is something much more problematic about issuing those orders to create such a person for the sake of his/her skin. Here, however, my hesitation is not based in prudence but morality. From the point of view of narrow self-interest, creating such a person for his/her skin seems preferable even

[17] Some may think that a crucial factor in this scenario is that my duplicate never gains consciousness. Suppose instead that, for whatever reason, the skin donor must gain (or regain) consciousness before he is killed the next night. To me it still seems prudent to make the orders if I can also ensure that the person to be killed does not know (or entirely fails to remember) that he will be sacrificed for his skin the next night. This might be ensured by sequestering the branches separately for the day and by using an effective memory-inhibiting drug.

to creating a near duplicate of myself, assuming compatible skin. What makes the interpersonal sacrifice objectionable, therefore, is precisely that the person to be sacrificed is not me but someone else entirely. This suggests that the idea of me splitting into multiple distinct persons can lead to distortions in moral thinking that are only corrected by rejecting the very intuition that is supposed to be providing us with moral guidance: the intuition that the other persons are also oneself.

I conclude that Singer has not identified any fundamental feature of morality in which he can convincingly defend a general principle of beneficence needed for the shockingly demanding conclusion he wishes to draw. Singer's principle of equal consideration of interests is itself too controversial, and Singer's proposed strategy for deliberation fails to provide any suitable criterion for arbitrating the competing concerns of different persons.

1.3 The disanalogy challenge

Singer's third strategy for defending his general principle of beneficence is a combination of tactics to resist the criticism that his principle mistakenly requires action in specific kinds of cases, and, in particular, those kinds of cases that would make compliance alarmingly difficult (as when relatively affluent individuals can, through small donations to aid agencies, rescue some of "the nearly 10 million young children dying every year from avoidable, poverty-related causes" (2009: 8-9)). The initial move is to ask the critic to specify what it is about such cases that make them relevantly different from analogous cases that Singer constructs, cases that typically involve children dying in shallow ponds. Once the reason for exception is specified, Singer

considers its moral significance. Singer dismisses some of the proposed reasons, such as the consideration of mere physical proximity to the imperiled person, for their prima facie insignificance.[18] Singer dismisses other reasons by observing that such considerations intuitively fail to counteract the obligation to act in analogous cases, which are often variants of a case in which a child is drowning in a shallow pond. So, in response to the proposed reason for exception that there are others who could also help the "10 million young children," Singer responds with a question: "Should I consider that I am less obliged to pull the drowning child out of the pond if on looking around I see other people, no further away than I am, who have also noticed the child but are doing nothing?" (1972: 233).

Singer's use of these analogies draws heated criticism. It is, according to John Kekes, "a rhetorical stratagem that misleads the uncritical and infuriates the critical" (2002: 506). The basic problem, as I see it, is that these appeals to analogous cases cannot fulfill their intended purpose of demonstrating that the proposed reason for exception lacks the relevant moral significance (by the critics' own intuitive judgments or otherwise). To suppose otherwise is to overlook or disregard what Frances Kamm calls the Principle of Contextual Interaction: a given feature may behave differently in different contexts (2007: 17).[19] So, although the presence of others intuitively fails to counteract the moral obligation to help in the modified pond case, we cannot thereby conclude that the same feature doesn't affect the moral obligation in other contexts. Even

[18] More specifically, Singer dismisses the moral significance of physical proximity as *until then* undisputed (2004: 11), by appealing to the impartiality claims discussed previously (1972: 232), or simply by the intuitive moral insignificance of the feature when considered abstractly.

[19] Shelly Kagan discusses the fallacy of overlooking what amounts to Kamm's principle in his "The Additive Fallacy" (1988).

if we can trust our intuitive judgments about our moral obligations in the analogous cases, such judgments cannot therefore be used to demonstrate the *general* moral insignificance of any particular feature.

When discussing his use of analogies, Singer concedes that they can be "pushed too far." He suggests that Shallow Pond works by "pulling on your heartstrings" through "focusing on a single child in need," contrasting this with "appeals to your reason" (2009: 17). Elsewhere, Singer explains these cases as an appeal to his readers' intuitive judgments. Since, however, Singer is willing to be somewhat dismissive of such judgments, he explains that for him "an appeal to intuitions is always an *ad hominem* argument" (Singer, Jamieson 1999: 315). Singer means that the appeal to intuitions is a device by which he "might win certain kinds of opponents over to his own views," views which Singer finds compelling on independent grounds (Singer 1974: 498).

Such responses do not engage the criticism that I have just raised. If the point of the analogies is simply to focus our attention on a single child in need, it is hard to see why Singer cannot rather focus our attention by describing in detail one of the children perishing overseas. It is hard to see why he should appeal to analogous cases. If, however,"[t]he point of the analogy is to force people to recognize an inconsistency in their moral convictions" (Singer 2007: 480), it remains unclear what that inconsistency is supposed to be. Given the Principle of Contextual Interaction, the judgment that a given consideration fails to counteract the obligation in a case involving nearby drowning children is perfectly consistent with the judgment that the same consideration generally does create a reason for exception concerning our opportunities to rescue "the nearly 10 million young children dying every year from avoidable, poverty-related causes." I

conclude, therefore, that Singer is not actually attempting to give a response to the concern raised by the Principle of Contextual Interaction.

1.4 The empirical challenge

The fourth weakness concerns the empirical claims that are needed to show that a plausible general duty-specifying principle can generate demands with regard to the distant needy. These are concerns about the effectiveness of aid agencies, the unintended consequences of providing foreign aid, facts about how aid organizations actually work, and the best options individuals have for dealing with chronic poverty and its associated problems. Facts of these sorts threaten to undermine an empirical claim critical to the famine-relief argument, such as Singer's claim that, "[b]y donating to aid agencies, you can prevent suffering and death from lack of food, shelter, and medical care, without sacrificing anything nearly as important" (2009: 15-16).

2. Efforts to Improve the Famine-Relief Argument

Peter Unger and Garrett Cullity have endeavored to improve Singer's famine-relief argument. While these are not the only attempts to defend or improve the famine-relief argument, they are noteworthy in this literature for the sustained attention they give to the matter. I discuss these attempts with the goal of highlighting their limitations.

2.1 Unger's attempt

We can understand Unger's contribution as a response to the methodological challenge through a sustained attempt to undermine our intuitive resistance to the highly demanding conclusion. Unger tries to undermine this intuitive resistance by advancing a debunking hypothesis largely based on what he calls "informal moral psychology" (1996: 23). By examining our intuitive responses to numerous hypothetical cases, Unger believes he can identify the factors to which our intuitive judgments are sensitive. Since the operative factors cannot, according to Unger, plausibly sustain the normative significance our responses seem to give them, they are best viewed with suspicion. By explaining our intuitive resistance to Singer's highly demanding conclusion in terms of these psychologically significant factors, Unger tries to clear the way for us to regard the highly demanding conclusion as reflecting our "deepest moral commitments."

To make his case Unger tests our intuitive responses to hypothetical cases with the twin goals of dismissing common ideas about potentially normative factors that might seem to be driving our intuitive responses (e.g., "social distance," "informative directness," and "urgency") (33-54), and of inferring the factors that are influencing our intuitive responses (e.g., "fallacious futility thinking," "projective separating/grouping," and the conspicuousness of the need). Towards the first goal, Unger constructs imaginative cases that reveal, or are supposed to reveal, that the factors in question do not consistently alter the relevant intuitive judgment. Typically the compared cases vary in more than one way. As a result, to directly conclude that a given factor isn't psychologically relevant from our responses to such cases is to overlook a close cousin to the Principle of Contextual Interaction. That is, if the effect of the factor in question on our intuitive judgments can vary due to other features of the context, then we cannot

conclude that a given factor generally doesn't affect our intuitive responses from the mere fact that it fails to affect them in the particular cases that Unger provides.

Towards the second goal (that of identifying the factors that are influencing our intuitive responses to cases), Unger expands his investigation to cases far beyond those simply having to do with help or rescue. While Unger's hypotheses often seem to do a good job at explaining many of the otherwise mysterious responses we intuitively make, his hypotheses are justified only to the extent that they provide the best explanations of these responses. As such, the arguments are only as strong as our reasons for concluding that there are no equally good alternative explanations for our intuitive judgments. In her detail analysis of Unger's efforts, Frances Kamm's alternative interpretations of his cases make such a conclusion difficult to reach (2007: ch. 6).

Even if Unger's explanations of our intuitive responses to particular cases are correct, it is a further step to show that these explanations debunk our intuitive responses. We would still need to show that factors that Unger identifies as psychologically influential are normatively insignificant. This further task isn't easily accomplished. While the identified facts may, on their face, compellingly appear to be normatively insignificant, the further implications of that insignificance may prove to be more implausible than the denial of that compelling appearance. This was David Lewis's response to Unger's conclusions (2000: 157-58). Moreover, the apparent normative insignificance of a given factor may disappear once the factor is placed in the context of a richer explanation. In discussions of the so-called special duties, one commonly encounters explanations for the normative significance of a factor that, from certain angles, appears as if it should be without significance (such as the fact that one happens

to have been born to these parents rather than those, or that one was born on this side of the border rather than that side).

Finally, even if Unger succeeds both in identifying the psychological factors that are generating our intuitive responses to cases, and in showing the normative insignificance of such factors, this only shows that we should not trust a conclusion *because* of its correspondence to our intuitive responses about cases.[20] While such debunking explanations would give us grounds for heightened scrutiny of theories that might be biased to align with our intuitive responses, they would still not supply everything we need to reach Singer's conclusion. We may still have additional grounds for resisting Singer's conclusion. In particular, we may simply think that Singer and Unger have misinterpreted our deepest moral commitments or the relevant moral principles. We therefore still lack the positive guidance of the sort that the general duty-specifying principle is supposed to provide.

A lack of positive guidance gravely weakens Unger's debunking explanations of our intuitive resistance to the demanding conclusion. Even if we accept that our intuitive responses to cases are overly influenced by factors which, of themselves, are morally empty, we do not yet know, for any particular case, whether or not these factors are distorting or correcting our intuitive judgments. Unger, for example, sometimes claims that one morally empty factor is correcting the distorting influence of another morally empty factor, but these claims about which factor is distorting and which is correcting

[20] As Singer puts it, "[Unger's] explanations are devastating for the view that we should take our intuitive responses to particular cases as a test of a sound theory, because the explanations show that our intuitive judgments are based on things that are obviously of no moral significance at all" (Singer, Jamieson 1999: 316). And so, "Post-Unger, no-one should place weight on our intuitive responses to such cases until they have empirically tested, and refuted, Unger's account of them" (Singer 1999: 187).

depend on an independent method of determining the correct judgment we should make.[21]

In summary, while Unger has moved us beyond the methodological challenge per se, this is only to push us into the controversial thickets that lay just beyond it, where his own proposals face challenges on multiple fronts.

2.2 Cullity's attempt

Garrett Cullity devotes the second chapter of *The Moral Demands of Affluence* to improving Singer's famine-relief argument. He calls his new version the argument from beneficence. This argument begins with the scenario in which he "could easily save someone's life which is threatened right in front of [him], and there is no further countervailing consideration [against performing the rescue], beyond the small effort it would cost [him] to help" (2004: 19). Failing to perform the rescue in such a scenario would be morally wrong, and Cullity explains this wrongness as an absence of sufficient "discriminating concern for other people's interests," or, for short, a "failure of beneficence." Such failures occur, according to Cullity, whenever one fails to act when

[21] Unger claims that we have "two main guides" in theorizing about morality. The first guide is our moral intuitions about particular cases. Unger's arguments are for the most part aimed at showing how this guide tends to be influenced by factors that seemingly shouldn't carry much moral significance. The second guide is what Unger calls our "general moral common sense." This "sensibility" is supposed to be like the first except that it is "directed at matters at least somewhat more general": it is "[p]itched at a level somewhere between the extremely general considerations dominating the tenets of traditional moral theories, on the one hand, and the quite fine-grained ones often dominating the particular cases philosophers present, on the other." Unger suggests that this second guide is more likely to reflect our deepest moral commitments (our "Basic Moral Values") and, "less directly, morality itself" (1996: 28). What I find missing in Unger's arguments is a clear articulation and defense of the supposed content of this general moral common sense—clear enough, at least, to help us determine when a factor influencing our intuitive responses about cases is distorting, as opposed to correcting, our moral judgment. As my own proposals in chapters four and five will ultimately suggest, Unger neglects certain considerations that our "basic moral values" plausibly accommodate, in particular, the cumulative costs to an individual of compliance with a rule, alongside the distinctive values of localized fellowship and solidarity that would, in addition, be lost if our practice were to treat the perils of strangers nearby as dismissively as we are prone to treating the perils of strangers far away.

one's reason for helping another person is very strong and the opposing considerations aren't sufficiently significant. Cullity then argues that, by parity of reasoning, the failure to make life-saving donations to aid agencies is also an impermissible failure of beneficence.

We might object to Cullity's conclusion by arguing that differences between the two cases show either that the reasons for helping in both cases aren't equivalent, or that there are stronger countervailing considerations in the second case. Cullity believes he can resist these objections, and he illustrates his strategy by considering whether differences in the "immediacy" of the need might render the two cases sufficiently incomparable. An impermissible failure of beneficence occurs if an agent is "too self-regarding." To say that the immediacy of the need in the first case is part of the agent's strong reason for helping in that case would make the agent's reason too self-regarding (21-23). Likewise, to say that the non-immediacy of the need in the second case undermines the reason for helping would make the agent's justification for not helping too self-regarding (25). The consideration of immediacy is too self-regarding simply because it appeals to one's own proximity to the need and the reasons that are "characteristic of a beneficent person to recognize" do not contain self-reference (22-23).

Cullity's strategy is broadly applicable inasmuch as it treats as too self-regarding any consideration that contains reference to the potential benefactor. This is a remarkable standard. I can readily imagine various possible facts that contain reference to myself and which are nonetheless highly relevant to my reason for helping others in need: that I am too ill at the moment, that helping the person would require me to sacrifice my own life, that I am currently occupied with the task of defusing a terrorist's bomb. Moreover,

it seems misleading to say that "the force of the reason for me to help someone who needs it comes from a fact about him, not a fact about me." The force of my reason to help another person may, on the contrary, have quite a lot to do with the fact that I am, for example, in a particularly good position to offer help, or that helping the other person would cost me very little. Far from introducing too much self-regard, in some of these examples the self-reference introduces no self-regard, at least if we take self-regard to mean a regard for one's self-interest. Since reference to the self can have so little to do with self-regard, I think we must reject Cullity's criterion.

If instead of self-reference we focus on the idea of self-regard, we still need a way of showing that a given consideration contains too much self-regard. In discussing the beneficent agent's reason for rescuing someone, Cullity apparently thinks that any element of self-regard in such a reason is too much self-regard. Cullity can think this because of the way he differentiates the agent's reason for helping from the considerations that countervail against that reason. With this distinction, Cullity would resist the idea that his reason for helping the other person is that the other person needs help *and* it takes him (only) a minute to provide that help. Rather, Cullity would insist that his reason for helping the other is simply that the other person needs help. The fact that it takes him only a minute to provide that help factors not into the reason for helping, but rather into the "significance of any considerations that countervail against that reason" for helping. Generally, any consideration one appeals to in one's justification for not helping is separated off as a countervailing reason. This distinction explains Cullity's suggestion that, in one's reason for helping another, any amount of self-regard can be too

much self-regard. On some occasions any degree of self-regard in one's reason for helping another will imply insufficient beneficence.

I see no point in challenging Cullity's distinction between reasons for helping and countervailing reasons. We can instead focus on Cullity's standard for determining when a justification for not helping is too self-regarding. Cullity claims that an appeal to the non-immediacy of the need in justifying inaction would be too self-regarding because it implies that "the force of the reason for me to help" is weakened by "facts about me," rather than facts about the person in need (25). This response implies, implausibly, that facts about oneself—or, more narrowly, self-regarding considerations—can never make it permissible to refuse help to another person. However, strong self-regarding reasons (such as the fact that offering help would place my own life at risk) can plausibly weaken or even block my duty to help, even when there are strong considerations that count in favor of helping another person (perhaps, e.g., the person will suffer and soon die otherwise). Cullity's argument from beneficence fails to provide a plausible criterion for determining if a given consideration is too self-regarding to countervail against one's reason for rescuing others. His rejection of non-immediacy as a countervailing consideration is flawed and he provides no reason for thinking that there aren't other considerations that countervail against our reasons for making life-saving donations to aid-agencies. It is understandable why, in his otherwise favorable review of Cullity's book, Singer himself doesn't think the argument from beneficence makes a significant advance in their shared goal of creating an argument with broad-based appeal (Singer 2007: 476).

Cullity's other contribution to improving Singer's argument is his response to the empirical challenge, including his sustained efforts to address the fact that effects of individual contributions to aid agencies tend to be thinly dispersed, with the seemingly relevant consequence that individual contributions typically have no significant impact for any beneficiary (1996, 2000, 2004: chapters 3-4). While I don't want to minimize these contributions to the famine-relief argument as such, my peculiar interests in that argument will allow me to side-step these empirical controversies.

CHAPTER THREE:
THE DIFFERENTIATION QUESTION

In the last chapter I considered the limitations of both the famine-relief argument and two attempts to improve it. My present aim isn't to exploit these limitations to rebut the highly demanding conclusion, but rather to articulate the legitimate challenge that remains. The famine-relief argument still poses a pressing challenge even for those who do not share Singer's utilitarian assumptions, Unger's conclusions from informal moral psychology, or Cullity's empirical claims about aid organizations and their effectiveness. This chapter's articulation of the challenge sets the stage for a fuller investigation into what it is that makes this challenge so difficult to resolve.

1. THE DIFFERENTIATION QUESTION

Let us say that I *rescue* another person if and only if I cause someone or other to be saved from serious peril whose likelihood to escape such peril is otherwise extremely low (less than one percent, we can say). *Serious peril* is intense, debilitating suffering, which, barring intervention, is immediately followed by death. Even granting these

restricted definitions, I am not morally required to act in every occasion in which I have the opportunity to rescue someone, but sometimes I am. A paradigmatic case in which I am so required is

> *Shallow Pond*: on a leisurely walk, I find myself uniquely positioned to rescue a drowning toddler from a knee-deep pond. Performing the rescue would be easy and safe, but I will ruin the new shoes I bought only a few days ago. (Adapted from Singer 2009: 3)

A paradigmatic case in which, perhaps, I am not so required is

> *Distant Rescue*: I am well acquainted with an international aid organization, Heroes for Hamiltons, which identifies neglected opportunities for saving people from serious peril at very low costs. Since the organization always identifies far more potential beneficiaries than its budget allows it to rescue, every $10 donation it receives will be used for the rescue of some additional person. (At the time of donation, this person is randomly selected from the large and continually replenishing pool of candidates, most of whom will perish.) Heroes for Hamiltons can therefore guarantee that, for every $10 I donate, they will rescue someone or other. I can conveniently send such donations via mobile phone text message, and I know that the organization has an excellent track record.

In each of the above cases the *salient fact* obtains: I can rescue another person at a small personal cost. I call it the salient fact because, when considered in the abstract, it appears to possess considerable normative significance. If I am required to rescue others whenever the salient fact obtains, then we should draw the *demanding conclusion*: I must rescue others so long as the salient fact obtains. A *moderate* account of our moral requirements of rescue has two features: it rejects the demanding conclusion in Distant Rescue, and it affirms the duty to perform the rescue in cases like Shallow Pond.

A philosophically satisfying defense of the moderate position must explain why I should not regard Shallow Pond as presenting just an additional opportunity of rescue right alongside those presented in Distant Rescue. Otherwise, so long as ten dollars

remains a small cost to me, I ought to regard myself as obligated to donate to Heroes for Hamiltons whenever I am obligated to rescue the imperiled person in Shallow Pond.

To make the line of reasoning here more explicit, we can say that, for any number *n*, Distant Rescue$_n$ is the scenario in which I have already contributed enough money to rescue *n* imperiled people. Let *m* be the maximum number of ten-dollar donations I could contribute to Heroes for Hamiltons such that I am still in a state in which ten dollars constitutes a small cost to myself. *Distant Rescue + Shallow Pond* is the scenario in which the two cases obtain simultaneously. That is, it is the case in which I encounter the imperiled person in Shallow Pond while also faced with the opportunities to rescue the distant needy through Heroes for Hamiltons. The moderate must avoid this argument's conclusion:

1. For all $n < m$, I am obligated to rescue the nearby drowning person in Distant Rescue$_n$ + Shallow Pond.

2. For all $n < m$, if I am obligated to rescue the nearby drowning person in Distant Rescue$_n$ + Shallow Pond, then I am obligated to donate $10 in Distant Rescue$_n$.

3. Therefore, for all $n < m$, I am obligated to donate $10 in Distant Rescue$_n$.

The moderate must in particular find a way to reject the second premise. In so doing, the moderate will be rejecting the idea that the imperiled person in Shallow Pond should be regarded as just another person whom she might rescue at a small cost to herself. The central philosophical challenge facing the moderate is the challenge of justifying the intuitive moral difference between Shallow Pond and Distant Rescue. At the center of this challenge is a question: given the standing opportunity to rescue others at small cost through Heroes for Hamiltons, why, at some point, must I still rescue the person who

would perish for want of my action in Shallow Pond, but not the person who would perish for want of my action in Distant Rescue?

Plausible responses to this question will account for widely shared and stable intuitive judgments that affirm similar moral judgments about variants of the Shallow Pond and Distant Rescue. These two cases can be treated as paradigms of two kinds of rescue cases. *Pond-like cases* are those rescue scenarios that an intuitively plausible moderate position will treat like Shallow Pond. Generally, whatever other opportunities one may have to rescue others, pond-like cases are those that continue to generate a duty to perform a specific act of rescue. *Famine-like cases* are cases like Distant Rescue. Although they also present an opportunity to rescue another person at a small cost, they appear, at least to the non-idiosyncratic moderate, not to generate a similarly strict duty of rescue regardless of what one has already done. The names are meant to recall Peter Singer's examples (1972) without assuming or identifying any principled basis for treating the cases differently (it should be obvious that ponds and famines are inessential).

The *differentiation question*, then, is the central philosophical question posed by the famine-relief argument to the moderate. It asks the moderate why one must continue to perform rescues in pond-like cases but not in famine-like cases.[22]

[22] Unger also frames the question of accounting for the intuitive normative difference between two types of cases when he presents "Vintage Sedan" alongside "The Envelope" (1996: 24-25). He then attempts to make the challenge more forceful by his comparison of "Bob's Bugatti" and "Ray's Big Request" (135-37). In constructing his pairs, Unger tries to make the differences between the cases such that, when considered in the abstract, those differences should actually strengthen the reason for acting in the case in which the rescue is more like Distant Rescue and less like Shallow Pond (in the former case, e.g., more people will be rescued, the cost of rescue is far less, etc.). By appealing to our intuitive response to Bob's Bugatti (a case of in which the rescue of a nearby stranger is extremely costly to the rescuer), Unger argues (with the aid of the conclusions drawn from his efforts at informal moral psychology) that we are obligated to make similarly costly sacrifices to save the lives of distant strangers (1996: ch. 6).

2. Pressing the Differentiation Question

As a question of explaining an intuitive normative difference between two types of cases, the differentiation question is not unlike the trolley problem. Many of us intuitively discern what we take to be a difference in moral permissibility between turning the trolley so that it runs over the single person instead of the five, and sacrificing the one healthy patient so that her vital organs can be used to save the lives of five others. It has also, however, proven difficult to account for this difference.

While the differentiation question is like the trolley problem in these respects, I believe it is made more difficult by an additional source of skepticism. The moderate's resistance to the demanding conclusion in Distant Rescue is plausibly motivated by the moderate's own will. While it is, for many, difficult to give up the idea that turning the trolley is morally permissible (or the idea that sacrificing the one is a permissible means to saving the many), there is a further kind of difficulty that the moderate must overcome in accepting the demanding conclusion in the famine-relief cases. In the differentiation question, the moderate's own will and sense of moral rectitude are plausibly strengthening her resistance to that conclusion in a suspicious way. Elements of the famine-relief argument draw out these suspicions. Showing just how they do will reveal what is right about the famine-relief argument.

2.1 A general duty-specifying principle

As discussed in chapter two, all variants of the famine-relief argument appeal to a general duty-specifying principle, the support of which is characteristically problematic.

I will argue for a plausible version of such a principle that will serve the more modest task of making a presumptive case against the moderate, transforming the differentiation question from a puzzle into a pressing challenge.

In Shallow Pond and Distant Rescue, the salient fact obtains. That is, in each case I can rescue another person from serious peril at a small personal cost. When considered in the abstract, the salient fact seems to be morally significant. To affirm this significance is to affirm much less than some ideal of impartiality, universality, or of respecting others' interests or preferences as I do my own, or even at some discounted rate below my own. Rather, it is only to affirm that, when I can easily do something to relieve it, the intense, debilitating suffering and survival of another person as such has moral significance for me. In speaking of "another person *as such*" I mean to refer to persons as mere persons, as not as siblings, friends, or others to whom we bear some special relationship.

Rejecting such an affirmation would be difficult. To see this, we might set aside the occasions in which I am, or might through my decision become, causally responsible for the suffering and threat of death of the other person. If, in the remaining possible circumstances, the suffering and survival of another person as such isn't morally significant to me *when I can easily provide the remedy,* then, it is hard to think why such considerations should ever be morally significant to me. Moreover, if the *suffering and survival* of another person as such isn't morally significant to me, then it is hard think why any of the interests or concerns of another person as such should be morally significant to me. It is, however, hard to think that anything should be morally significant to me if the interests and concerns of another person as such are never are. So, if

anything is morally significant to me outside of special relationships and occasions in which I am or would be causally implicated in someone else's harm, then the salient fact plausibly is. Put another way, acceptance of any positive duty of beneficence towards other persons as such would be difficult to square with the denial of the salient fact's moral significance.

Accepting that there is some general duty of beneficence towards other persons as such, the salient fact should not only be significant, but highly significant, given the particular kind of opportunity it presents. It presents not simply an opportunity to do better for someone else than I can otherwise do for myself, but rather an opportunity to do *much* better for someone. Moreover, it is not simply that there is a great disparity between what is at stake for me and what is at stake for the other person whom I might affect (as when I could rescue a person from death by sacrificing one of my arms; or when I could save a person from walking a mile by giving a ride). Rather, what is at stake for the other person concerns rescue from serious peril, and what is at stake for me is, in contrast, a very small cost.

When we attempt to specify the intuitive significance of the salient fact through a general moral principle defining our duties, we encounter a difficulty common to all such tasks: for whatever duty-defining principle we specify, there almost always seem to be exceptions. So, although the salient fact seems to be the sort of fact that might give rise to a moral obligation, we can easily think of cases in which the salient fact doesn't suffice to create a moral obligation (e.g., when the rescue would endanger the lives of others, when it would sacrifice opportunities to do even better things, when it would encourage others to get themselves into similarly serious peril, etc.). The salient fact, therefore,

doesn't always give rise to a moral obligation to rescue the person. It might even be that, when it obtains, the salient fact doesn't even usually give rise to the moral obligation. Perhaps there are usually further considerations that effectively counteract any such obligation.

We can acknowledge the intuitive significance of the salient fact while still respecting these exceptions by saying that the salient fact is morally significant by default. More specifically, we can say that the salient fact obligates me to save the person from serious peril so long as that duty isn't effectively counteracted by other features of my circumstance. To say, in the intended sense, that a consideration[23] *counteracts* the presumptive duty is to say that it undercuts or overrides the presumptive duty to save the person (it is not enough if the consideration merely provides opposing normative force, or that it to some extent weakens the reasons to perform the duty). Here, then, is a plausible general duty-specifying principle:

> *Rescue principle*: If I can rescue someone from serious peril at a small cost to myself, and there are no counteracting considerations, then I am morally required to rescue the person.[24]

The rescue principle says that, by default, I have a duty to perform the rescue when the salient fact obtains. Our recognition of this status is, I think, reflected in the presumptions that would be associated with the salient fact. To see this, suppose my new

[23] It may be that what counteracts the presumptive duty to perform the rescue is most easily thought of not as a single consideration, but rather a set of considerations, or of one consideration interacting with others. Since, however, we may speak of a set of considerations as itself "a consideration," I will not always make the explicit distinction between single considerations and sets thereof.

[24] I am here borrowing from Garrett Cullity (2004: 17-19). Many versions of the general duty-specifying principle have been offered, but the one I settle on is perhaps closest to Cullity's version: "[W]hen someone's interests will be very seriously compromised if I do not help him, but there is no seriously countervailing consideration, then not helping is…morally wrong" (19). I prefer to use "counteract" instead of "countervail" since the latter term more readily suggests considerations that override instead of undercut. Cullity, however, is careful to define "countervail" so as to include considerations of both kinds.

neighbor learns the following random fact about me: I, his new neighbor, have recently refused an opportunity to easily rescue another person from serious suffering and death. My sense is that such a fact is bound to reflect poorly upon me. While my new neighbor might recognize that my behavior may have been excusable or even entirely appropriate (given other special considerations undisclosed), the bald fact seems to create a presumption to the contrary.[25]

The rescue principle gains plausibility from its modesty. It only claims that the salient fact creates a moral duty if other factors don't create exceptions. As to what these exceptions might be, the door is left open. Left open is the possibility that, because the salient fact doesn't have much force or significance, these exceptions are easy to come by. Given its accommodation of such interpretations, the rescue principle is difficult to deny. While theories about morality will differ in how they explain the rescue principle's legitimacy, I expect that most sensible theories will accept it.[26]

Despite its ability to accommodate wide interpretations of the salient fact's significance, the rescue principle still carries force. The principle claims much more than simply that we have duties to rescue imperiled persons except when we don't. The rescue principle resembles any moral principle, except that it makes explicit the room for exceptions that we tend to tacitly recognize in other principles. For example: you ought not to intentionally deceive another person *at least in the absence of counteracting*

[25] It might be thought that the neighbor's reaction is likely to be influenced by various types of implicatures attached to how he has received the information. If a person conveys such information to my neighbor it is likely to be interpreted as a warning, implying that there were no special counteracting reasons for my behavior. So we might imagine that my neighbor comes by this information about me through a device that simply spits out random facts, most of which are entirely pointless. In such a case, the bald fact still seems to reflect poorly upon me by default.

[26] This provides the sense in which a demandingness problem grounded in the rescue principle differs from the second category of demandingness problems discussed in chapter one, where the demandingness requires the assumption of a principle peculiar to a more specific theory.

considerations—which sometimes obtain. We cannot, moreover, build a principle out of just any fact. The salient fact is special. We would not, for example, say that I am, by default, morally required to kick off my shoes whenever someone drops a penny. The fact that someone has dropped a penny doesn't even count in favor of kicking off my shoes; there is simply nothing to counteract. The rescue principle therefore attributes a special kind of moral significance in the salient fact, and this is a status it shares with the kinds of considerations generally found in the principles that state our moral duties.

We might better understand this special kind of significance that the rescue principle gives to the salient fact by considering what a counteracting consideration would have to be like. In Shallow Pond, the salient fact obtains. I assume, with the moderate, that it is morally wrong not to save the imperiled child in Shallow Pond. More tentatively, I find it plausible that this wrongness is closely associated with the salient fact, and, in particular, that the refusal to rescue the child is wrong because it violates the rescue principle.

2.2 Appeals to analogy

Despite the plausibility of the rescue principle, and its applicability to Shallow Pond, it may be that the wrongness of refusing to save the drowning child results from some other features of that scenario. Or it may be that the wrongness of such a refusal is over-determined in Shallow Pond. That is, there may be other features of Shallow Pond that would also suffice to require one to rescue the drowning child. Finally, it may be that certain features of Shallow Pond function to block considerations that would, in other circumstances, tend to counteract the duty to save an imperiled person. We mustn't

assume that a consideration that is morally significant under one set of circumstances carries the same moral significance under different circumstances.

We should remember these possibilities whenever we are tempted to use cases like Shallow Pond to draw conclusions about the suitability of a proposed counteracting consideration for the rescue principle. As discussed in chapter two, this is a common, yet questionable, strategy of proponents of the famine-relief argument. When it is suggested that a given consideration effectively counteracts the obligation to save a person in a famine-like case, Shallow Pond is then modified to incorporate that consideration. Someone might suggest that, in some situations in which the salient fact obtains, the presence of others who are better positioned to save the imperiled person would counteract my obligation to perform the rescue. If, however, we populate a pond-like case with a couple of other passersby who are better positioned to help the drowning child, but who are neglecting to do so, we intuitively judge that the proposed consideration does not counteract the moral duty to save the child in the shallow pond case. I will call this argumentative move an *appeal to analogy*.

2.3 Appeals to analogy and skepticism of the will

While in the last chapter I pointed out the limitations of appeals to analogy, here I want to argue that there is a way in which such appeals increase the forcefulness of the differentiation question through an appropriate kind of skepticism towards the moderate. The appeals to analogy strike many people as intuitively forceful and I want to identify a sense in which such intuitions are right.

The popularity of appeals to analogy would be explained by a popular assumption that a given factor makes the same difference everywhere if it makes that difference anywhere. Shelly Kagan argues that this "ubiquity thesis," as he calls it, is supported by the attractive but dubious assumptions that (a) "the status of the act is the net balance or sum which is the result of adding up the separate positive and negative effects of the individual factors" (the "additive assumption"), and that (b) "variations in the other factors will not affect the contribution made by the factor in question" (this is the denial of Kamm's Principle of Contextual Interaction, mentioned previously) (Kagan 1988). Since these assumptions are so questionable, Kagan's explanation would seem to debunk the appeals to analogy. This debunking explanation seems particularly plausible when appeals to analogy are offered as demonstrations of the general moral insignificance of a given factor.

It is tempting, therefore, to reject all of the appeals analogies as sophistry aimed at persuading critics of the famine-relief argument to concede more than they really should. However, there is also a non-debunking explanation. We can alternatively view such appeals as exploiting a plausible extension to the universality requirement of reasons, or of judgments about reasons. A plausible version of this requirement of reason finds universality in the very concept of a reason: "the claim that *X* is a reason for one person to do *A* entails the claim that it is also a reason for anyone else whose circumstances are similar in relevant respects" (Scanlon 2003c: 177-8). Similarly, the claim that some factor counteracts someone's duty to perform a rescue might plausibly be thought to entail the claim that the same factor would counteract the duty of anyone else whose circumstances are similar in relevant respects. The claim that a given factor counteracts

the presumptive duty in one case in which the salient fact obtains entails the claim that the same factor counteracts the presumptive duty across the range of relevantly similar cases.

Insofar as this extension maintains the plausibility of the universality requirement of reasons, it shows very little. In particular, it doesn't show that a proposed countervailing consideration in Distant Rescue must also counteract the duty in Shallow Pond. It is entirely possible that the pond-like cases typically aren't "relevantly similar" to the famine-like cases. So, even though we intuitively judge that the presence of others in a better position to help does not counteract the duty to rescue the drowning child in Shallow Pond, it can always be denied that Shallow Pond is relevantly similar to Distant Rescue.

Such a denial of relevant similarity, however, shouldn't be automatically granted. In particular, if we insist that the pond-like cases aren't relevantly similar to the famine-like cases, there is some pressure to say why this is so. If we cannot explain why the cases aren't relevantly similar, then the appeals to analogy reveal that we lack a principled way of determining when the proffered consideration counteracts a presumptive duty and when it doesn't. While this kind of inability may not be worrisome in some contexts, I think that it makes our attempts to offer counteracting considerations in the face of the salient fact worrisomely ad hoc.

Ad hoc proposals of counteracting considerations are worrisome in the face of the salient fact for two reasons. First, such counteracting considerations are supposed to counteract a feature that is, when directly considered as described, intuitively significant: one could rescue someone from serious peril at a small cost to oneself. Second, these

counteracting considerations create permissions that typically favor those who propose them. As the differentiation question shows, without some such proposal for limiting obligations in the face of the salient fact, the demanding conclusion follows. That conclusion, as I will later argue, means that morality threatens to make demands that many of us would be unwilling to meet. So, although a counteracting consideration must only counteract the salient fact in relevantly similar circumstances, I am suggesting that appeals to analogy can sometimes show that our denials of relevant similarity are ad hoc, and worrisomely so given the intuitive significance of the salient fact and the moderate suspicion we should have towards proposals that advantage ourselves. To answer in a satisfying way the concerns raised by appeals to analogy, we need to specify a counteracting consideration that answers to the intuitive significance of the salient fact, that avoids licensing morally wrong inaction, and that does not rely on an ad hoc application to certain cases but not to others.

The difficulty of finding a compelling counteracting consideration is well attested in the literature that addresses the famine-relief argument, some of which I will consider in the next chapter. Many differences separate pond-like cases from famine-like cases, and many counteracting factors have been proposed. Still, even if no single proposal is of itself satisfactorily plausible, the moderate might take comfort in the thought that some suitable counteracting consideration, or set thereof, is bound to exist. This insistence, however, would be awkward. The awkwardness would be similar to that of insisting that I have an excuse for my behavior, but that I just can't figure out what that excuse is. If the moderate is unable to specify the counteracting considerations, then respect for the

significance of the salient fact and a modest suspicion of motive should provoke some moral discomfort.

2.4 Side-stepping the empirical challenge

The denial of the demanding conclusion serves the interests of the one who occupies the position of the prospective beneficiary in Distant Rescue. This adds to our suspicions of the moderate's motives, which in turn pressures the moderate to answer the differentiation question. Distant Rescue, however, is a merely hypothetical scenario. Critics of humanitarian aid efforts cast doubts upon their effectiveness and raise concerns about their unintended consequences. It is difficult to believe, moreover, that each small individual contribution serves to rescue someone or other. Mindful of such things, the moderate may not feel threatened by the realization of a case like Distant Rescue when resisting the demanding conclusion. Skepticism towards that resistance may therefore be undermined by the purely hypothetical status of Distant Rescue.

The moderate's motivation to resist the demanding conclusion, however, doesn't depend on conceding the likelihood of Distant Rescue. The facts separating the supposed actual world from a world in which a famine-like case occurs are the sorts of facts that might change, or which we can readily imagine being different. The moderate can imagine discovering that Distant Rescue is far likelier than was formerly thought, or that the relevant facts have suddenly changed. It is easy to imagine, for example, discovering that there are many millions of seriously imperiled people overseas. We can further imagine that the Gates Foundation unveils tomorrow the creation of an organization of Heroes for Hamiltons to meet the needs of these people.

The first question to ask is whether these changes would so dramatically change the moral facts about what individuals are required to do. That I don't think so doesn't, at least to me, seem to be determined by my assessment of whether or not the hypothetical changes are a live possibility. More relevant, it seems, is my sense that the supposed changes in the empirical facts are not the sorts of changes that should effect such a dramatic change in my moral duties. My thought seems to be that, if I am not obligated now, then I shouldn't suddenly be so obligated after the mere sort of change effected by the Gates Foundation.

The mere conceptual possibility that my moral duties could be so dramatically altered by these sorts of changes is disquieting. If one's moral situation is so dramatically altered by *such* considerations, then acknowledging this may perhaps be enough to make one lose one's sense of security against such threats of demandingness generally. One's moral situation may seem too vulnerable to the wrong kinds of contingencies, even if the particular hypothetical contingency we have used to show this isn't likely to obtain.

Moreover, even if I think that I am not presently confronted by any famine-like case, I can make plans for what I would do if I were, or if I should suddenly learn that I am continually confronted by the salient fact. The force of these considerations about planning is clearest when I engage in first-personal inquiry. Let me therefore present a line of such inquiry that I expect will be shareable by others who are both attracted to the moderate position and earnest about the demands of morality.

I first realize that I can plan out what to do if confronted by Distant Rescue. Earnest about the demands of morality, I don't want to plan to do that which I believe to be, or even suspect to be, morally wrong, since to plan to do what is morally wrong is a

particularly bad kind of moral failing. On the one hand, it is unlike those failures that I can relegate to my past, and of which, as an earnest moral agent, I can come to terms with through remorse, a new commitment to doing what is right, and perhaps some compensatory behavior. On the other hand, it is also a failure that belongs to the merely foreseeable future, for which I might say with a shrug, "Try as I may, I sometimes fail." When I refuse to intend to fulfill my moral duties, this is a present refusal to adjust my inappropriate attitudes. To concede such a failure is to actively perpetuate it.

So, if I accept the demanding conclusion in Distant Rescue, then I must either refuse to plan what to do should such a case arise, or I must plan to comply with the demanding conclusion in that case. When I consider doing the latter, I discover it to be very difficult. I find it difficult to say, as a statement of intention, that I will comply with the demanding conclusion should certain contingent facts suddenly change.

Possibly, my resistance to forming such a plan is appropriate. I may resist such a plan because I doubt that the demanding conclusion in Distant Rescue is correct. That is, my doubts about the demanding conclusion in such a case may explain my resistance to it. Similarly, I would resist the demand to mow my neighbor's lawn should he leave on a vacation because I doubt the legitimacy of such a demand. However, the order of explanation may be the reverse. I might resist accepting the demanding conclusion in Distant Rescue because of my difficulty in adopting a plan to comply with it should such a case arise.

The latter explanation is disconcertingly plausible. In the lawn-mowing case, I can imagine the conditions under which I would change my mind and adopt the plan to mow my neighbor's lawn. Often, when I am unwilling to do something, I can identify

conditions under which I would become willing. This doesn't happen when I think about the demanding conclusion, where my unwillingness to comply is remarkably robust. That unwillingness persists under a wide variety of conditions, and under a wide variety of cases in which the salient fact persistently obtains. Whenever the salient fact persistently obtains, it is difficult for me to specify any further conditions under which I would be willing to plan to comply with the demanding conclusion. This suggests either that my grounds for doubting the demanding conclusion are robust in the same way, or that it is deeply entrenched unwillingness to comply with the demanding conclusion under that broad range of cases that is driving my doubts about that conclusion.

To show that it isn't my own unwillingness to comply with the demanding conclusion that is driving my resistance to it under such a broad range of cases, I might try to provide a case in which I would be willing to comply with a similarly demanding conclusion on behalf of the needs of strangers. Failing that, I might try to show that I have grounds for doubting that conclusion under the broad range of cases in which I resist the demanding conclusion. Whether or not I can do this will depend in part on the breadth of that range. These grounds will have to justify my resistance to the demanding conclusion not only in Distant Rescue, but also in all of the other variants in which I likewise resist the conclusion. So here again, appeals to cases that are in this way analogous to Distant Rescue provide legitimate challenges.

The challenge of planning is one that I place upon myself. (Pitching it to others provides a sense to Singer's idea that his famine-relief argument presents an *ad hominem* challenge.) I might think that I can evade the challenge by simply refusing to plan what I will do if famine-like cases should arise. If, however, my refusal to so plan is motivated

by my desire to avoid this challenge, then I can no longer say that I am making a good-faith inquiry into my attraction to the moderate position.[27] Inasmuch as making this good-faith inquiry is a part of what it means to be earnest about the demands of morality in approaching the differentiation question, this is not an option I can take. I must instead answer the differentiation question in a way that reveals a justification for refusing to form the intention to comply with the demanding conclusion not only in Distant Rescue, but also in a range of analogous cases where I similarly resist forming an intention to comply. Failing this will make it all too plausible that my attraction to the moderate position is better explained by a fundamental unwillingness to comply with the kinds of demands that the rescue principle might generate.

2.5 Taking stock

In this chapter I have extracted a challenge posed by the famine-relief argument. This challenge is the task facing the moderate of providing a plausible response to the differentiation question. My argument for why this task is particularly pressing for the moderate follows the pattern of Singer's own argument without crossing the limitations of that argument identified in chapter two. I proposed a general duty-specifying principle, but one whose plausibility does not depend on meeting the grounding challenge through a defense of the sorts of purported fundamental features of morality that Singer's principles require. Following Singer, I also used appeals to analogy, but I avoided the

[27] I might alternatively say that my refusal to plan what to do in Distant Rescue is motivated by the difficulty of determining whether or not the demanding conclusion is correct in such a case. Such a difficulty, however, arises because it is difficult to either affirm or reject the demanding conclusion in such a case. So, the question just shifts to why I think it is difficult to either affirm or reject the demanding conclusion in such cases. The debunking explanation of why I think this has just as much force so long as I cannot better explain my resistance to affirming the demanding conclusion in such cases. As previously discussed, I cannot escape by simply saying that I might have good grounds for resisting the demanding conclusion in such cases, despite my inability to identify those grounds.

disanalogy challenge by giving such appeals a role that avoids violation of the Principle of Contextual Interaction. I implicitly addressed the methodological challenge by raising skeptical concerns I raised about the moderate's motivations. Such skeptical concerns target the moderate's resistance to the demanding conclusion, suggesting a way to debunk the counterintuitive force of that conclusion. Finally, I side-stepped the empirical challenge by treating the famine-like cases as purely hypothetical, suggesting how their mere proximity to the actual world can still plausibly motivate the moderate's position.

My results, of course, are more modest than those sought by either Singer or Unger. I don't draw any shockingly demanding conclusion about what we as individuals are morally required to do presently. I don't even draw the demanding conclusion for the hypothetical famine-like cases. I have rather identified what I take to be the central philosophical challenge presented by the famine-relief argument.

3. Are We Committing Some Fallacy?

In presenting and pressing this philosophical challenge there is room for suspicion that I have fostered some kind of fallacy, rendering foolish my further efforts to answer that question. A prominent worry is that there is something deceptive in the way that I have treated our opportunities to rescue others in the famine-like cases. I have treated them as so many individual cases in which one can rescue imperiled individuals, by individual actions, each of which requires only a small sacrifice. I may be inducing error when I regard these opportunities as discrete and sequential events. Three versions of this worry are worth considering.

3.1 Iterating what we should aggregate?

The first version of the worry concerns Garrett Cullity's distinction between iteration and aggregation. To introduce this distinction, consider how David Lewis describes how we arrive at the shockingly demanding conclusion:

> An argument that is cogent once is cogent twice over. If indeed it is seriously wrong not to save the life of one distant child…then why is it not equally wrong not to save the life of the next distant child? And the next, and the next…? There is nothing to shut the argument off after you have saved one life. Or after you have sent [the funds] to save many lives….When you have so little left that it becomes doubtful whether you can live to give again another day, then the argument shuts off. But only then. (2000: 154)

Cullity sees the "iterative approach" in this form of reasoning:

> The iterative approach uses the following method to derive a conclusion concerning what an agent is morally required to do when presented with more than one person in extreme need: begin with what you would have been required to do for a single needy person, then iterate this for every other. (1996: 61-2; compare 1995: 194)

By "the iterative approach," Cullity seems to have in mind two distinct ideas. First, he seems to mean an approach whereby the question of whether or not one is morally required to act on behalf of an imperiled person is determined solely by the costs to the rescuer of helping that particular person. If those costs are really low, then you are presumably required to help. So, even when there are numerous such opportunities to help, you will presumably be required to help each time until you are so impoverished that giving even a few dollars would be, of itself, a significant sacrifice to you. What gets left out on this approach, or what is assumed to be insignificant, is any consideration of the cumulative costs of helping many such persons, including the sacrifices you have

already made, or of the sacrifices you might have to make in compliance with one's duties so assessed.[28]

So understood, the iterative approach supports the demanding conclusion in a dubious way. Such an approach simply assumes that the cumulative costs of your acting similarly (or even of your having already acted similarly) do not bear on the question of whether or not your are required to act so on a particular occasion—except if past costs have impoverished you to the point where an otherwise small additional cost becomes sufficiently dear to you. This is to disregard without argument sources of potentially counteracting considerations, considerations that may well be at least as forceful as any others. David Schmidtz rightly criticizes those who adopt this approach as considering only "Token-Costs" while negligently disregarding "Type-Costs" (2000: 691-93).

On a second interpretation, the iterative approach is simply the denial of the "aggregative approach," and, strangely, *the* alternative to it.[29] According to Cullity, the

[28] This interpretation is clearest when Cullity speaks of the iterative approach as one in which "[n]ot making [another contribution] will be wrong provided only that, considered in isolation (that is, apart from any contributions I have already made), not making it remains comparable to the direct and unextenuated failure to save a life" (2003: 413), and when Brad Hooker describes it like this: "If personal sacrifice is assessed iteratively, then a rule requiring you to help others in great need when this involves modest personal sacrifice would require you to help another in great need whenever the sacrifice involved in helping on that particular occasion is relatively little" (Hooker 1999: 180). Cullity also describes the approach as "holding that, for each needy person, it would be wrong for me not to help unless the cost of helping *that person* is large" (Cullity 1996: 59), or to involve "treating each person I could contribute towards helping as though he were the only person needing help" (2004: 173).

[29] Cullity writes, "The issue is not whether the cost to the agent of helping is an appropriately countervailing consideration in relation to beneficence. It is whether we should think of the cost to the agent iteratively or aggregatively" (2004: 87); also: "The alternative [to the iterative approach] is to take what I call the 'aggregative approach.' This denies what the iterative approach asserts…." (2003: 413). More significantly, the idea that these two approaches are our only (plausible) choices is also implicit in Cullity's form of argument: since the iterative approach leads to absurdity, we are justified in adopting the aggregative approach instead: "Part I has defended the life-saving analogy, but has shown that an iterative approach to it yields the Extreme Demand. Part II has shown that the Extreme Demand must be rejected. Together, then, they give us an argument for adopting the other possible approach to the life-saving analogy: an aggregative approach" (2004: 173). (My observations about the two distinct senses of the iterative approach therefore undermine Cullity's central argument. Even if he can show that the iterative approach—under the first interpretation—leads to absurdity, this doesn't justify that we should instead take

aggregative approach instructs us to "begin instead by assessing the magnitude of the overall collective need of the people who could be helped, then ask directly what overall sacrifice can be morally demanded of you in response to that collective need" (1996: 61-62); on this approach, "we should ask not about the size of the sacrifice I need to make to save each person, but about the size of the aggregate sacrifice that would be involved in helping them collectively" (2003: 413). Significantly, Cullity interprets the aggregative approach as "suggesting, in effect, that we should lump together all of the people I could contribute to helping, and ask what I should be prepared to do for this collective entity" (2004: 86).

Since, on the second interpretation, the iterative approach is simply the denial of the aggregative approach, it is clearly a mistake to equate the two interpretations of the iterative approach. To equate them would mean that the only alternative to thinking only in terms of overall needs and overall sacrifices is to consider only the costs of helping one particular person on one particular occasion. In other words, I must choose between thinking as if the person whom I might help is the only person I might help, and disregarding the fact that the cost to me of rescuing another person remains very small. There are, however, two potentially relevant considerations here and there is no obvious reason why I cannot keep my eye on them both. Put generally, the two considerations are (a) the cost to me of helping on the one occasion, and (b) the accumulating (or accumulated) cost to me of acting similarly (or of having acted similarly) on similar occasions. I see no reason why we cannot bear in mind both the fact that I can rescue another individual at a small cost to myself while nevertheless remaining open to the

---

the aggregative approach to assessing costs/duties, especially when the aggregative approach is characterized so distinctively.)

possibility that the cumulative costs of rescuing others (or of having had rescued others) can directly weigh into the question of how I am required to respond.

Keeping in mind both considerations is exactly what my framing of the challenge allows us to do. The rescue principle would have us keep our eyes on the salient fact (that I can rescue another individual at a small cost to myself), while remaining open to whether the accumulating costs of similar sacrifices figure in as (or to) a counteracting consideration.

Of course, although we should leave it as a possibility that facts about what I have already sacrificed, or by the expected cumulative costs to me of acting similarly on similar occasions, counteract the obligation to respond to the salient fact, we would want an argument for this. It isn't easy to see how such an argument might go. Such an argument would need to explain how that fact about the accumulating costs could carry sufficient force. Cullity expresses the difficulty this way:

> Why should my failure to save the hundred-and-first, or the thousand-and-first person be any more excusable than my failure to save the first, if the cost remains trifling compared to what is at stake for that person? My having already saved a thousand lives does nothing to alter the fact that I could save *this* person at an insignificant cost. (1996: 565)

The task of answering this question remains hard even though it doesn't yet force us to say, with Lewis, that "[t]here is nothing to shut the argument off after you have saved one life. Or after you have sent [the funds] to save many lives." Therefore, we can affirm the legitimate challenge without assuming the dubious version of the iterative approach.

3.2 A Sorites paradox?

The second version of the worry is the suspicion that the challenge prompts us to engage in a form of reasoning that generates a sort of Sorites paradox.[30] The Sorites paradox arises in cases in which a given property and its negation are connected by a continuum of incremental differences, each of which is seemingly trivial (e.g., the difference of losing a single hair), and, as a consequence of their triviality, no particular one of these incremental differences appears able to sustain the distinction between the property and its negation. As a result, there is no point at which we can plausibly say the property arises. Similarly, by conceiving of what we might be required to do on behalf of others on a continuum of small increments (from giving nothing to giving nearly all that I possess, in $10 increments), it should be similarly difficult to say precisely where we should draw the line between obligation and supererogation. Wherever we try to draw that line, we might always ask, "Why not $10 more, or less?" It is understandable, then, to think that, in thinking of the problem in this way, we have set ourselves an impossible task.

It is not only the size of the required sacrifice that can be placed on a paradox-inducing continuum. We might also be able to think of the two kinds of rescue cases (Shallow Pond and Distant Rescue) as standing on a continuum separated by many incremental and seemingly trivial differences.[31] In this way, then, the challenge of saying what exactly is supposed to constitute the morally significant difference between the two cases might seem like a task for which there can be no satisfactory answer.

---

[30] The Sorites paradox is discussed in relation to the famine-relief argument in Fishkin 1982: 54-59, Feinberg 1984: 64-65, and Cullity 2004: 176-77. Mothersill attributes this kind of reasoning to Shelly Kagan in his arguments for demanding conclusions (1993: 542-43).

[31] Kamm sees this strategy at work in Unger's arguments (2007: 199).

The mistake here, however, is to assume that the answer we are expecting is to be given in terms of a precise dollar amount, or in terms of an exact number of individuals we are required to save, or in terms of a single simple feature grounding the moral differences between the two types of rescue cases. I make no such assumptions. The challenge of identifying a counteracting consideration, or set thereof, that will block the rescue principle from becoming shockingly demanding does not require that the counteracting consideration be any simple feature distinguishing Distant Rescue from Shallow Pond, nor does it require us to specify a precise point at which we no longer have to go on saving the lives of distant strangers. Where the duty ends and when a consideration comes to counteract the presumptive duty may be vague matters.

If, moreover, in Distant Rescue there is pressure against drawing the line of duty at any precise number of lives to be saved or dollars to be donated, it has at least as much to do with the fact that giving just a little bit more money would rescue another person's life. Unlike the typical cases of the Sorites paradox, what counts against drawing a line is not simply the apparent arbitrariness of drawing the line in one place rather than another, whether the difference between these two alternatives is so seemingly trivial, but that an additional life can be saved if one gives just $10 more.

3.3 Like the paradox of the self-torturer?

The third related worry is that the form of the challenge is structurally similar to the conundrum facing Warren Quinn's self-torturer. Here is Quinn's scenario.

> *The Self-Torturer*:[32] Suppose that there is a medical device that enables doctors to apply electric current to the body in increments so tiny that the patient cannot feel them. The device has 1001 settings: 0 (off) and 1 … 1000. Suppose someone (call him the self-torturer) agrees to have the device, in some conveniently portable form, attached to him in return for the following conditions: The device is initially set at 0. At the start of each week he is allowed a period of free experimentation in which he may try out and compare different settings, after which the dial is returned to its previous position. At any other time, he has only two options—to stay put or to advance the dial one setting. But he may advance only one step each week, and he may *never* retreat. *At each advance he gets $10,000*....The trouble is that there *are* noticeable differences in comfort between settings that are sufficiently far apart. Indeed, if he keeps advancing, he can see that he will eventually reach settings that will be so painful that he would then gladly relinquish his fortune and return to 0. (Quinn 1990)

The self-torturer's predicament is that he can foresee a series of decisions which, when considered one-by-one, will foreseeably lead to an outcome that the self-torturer has very strong reason to avoid. The difficulty, however, is that there nevertheless appear to be good reasons to consider each new decision individually. This predicament seems similar to our own decisions when facing Distant Rescue. Although I foresee that the cumulative consequence of my continued contributions will be shockingly burdensome, there nevertheless appear to be quite compelling reasons to donate each additional $10. The difficulty here, moreover, clearly cannot be attributed to the iterative approach described above. Each individual decision still seems to be appropriate despite the self-torturer's continuous recognition of their cumulative costs.

If the self-torturer doesn't stop himself somewhere before the end of the series, he will end up having done something that is against his own self-interest. This is peculiar because every individual decision he would have made up to that point was seemingly in his own self-interest. The suggested parallel might be this: if, in $10 increments, I end up

[32] While I borrow the term "torturer" here, it is arguably misplaced. Torture doesn't just involve the infliction of pain, but also a specific relationship between the torturer and the tortured—which doesn't obtain in the case described.

giving all of the money that I can to Heroes for Hamiltons, then I will have ended up doing something that is well beyond what is morally required; this is peculiar because each particular donation seemed to have been morally required.

I have two reservations about this suggested analogy. First, although it is apparent that the self-torturer would have cumulatively done something against his own self-interest, it is not as apparent that I would have done beyond what is morally required of me in giving away all of my money to Heroes for Hamiltons. If I can be morally required to save one person for $10, then perhaps I really can be morally required to save *n* persons for $(10 x *n*), at least up to the point at which an additional $10 becomes extremely dear to me.

To circumvent this first reservation, it might therefore be helpful to merge the two scenarios, so that the cumulative cost to me is not primarily financial. Suppose that I wear a device very similar to that of the self-torturer, and that only by the click of the dial I can rescue another person (we should also stipulate that this is the only way these people can be saved and that this is the only benefit I can get from turning the dial). Am I morally required to "torture" myself for the sake of saving 1000 imperiled persons? Suppose that the difference between no electrical shock and a continually "tortuous" electrical shock can be bridged by only 20 increments, where the difference between adjacent increments to me is still trivial.[33] Can it be that I am morally required to inflict continuous agonizing pain upon myself for the sake of saving 20 imperiled people? Since I find it much easier to judge that I would not be morally required, we here have a

[33] There is nothing conceptually impossible about this, or at least not obviously. Larry Temkin (1986) considers the case of color perception. Borrowing that comparison, we might notice that relatively few paint chips may be needed to bridge the gap between a lovely shade of green and a truly awful one such that any perceived difference between any two adjacent chips is trivial. This feature of color perception, of course, might not be shared in the experience of pain.

case that might provide instructive similarities connecting the Self-Torturer to Heroes for Hamiltons. Call this new case the *Beneficent Self-Torturer*.

My second reservation about the suggested parallel is that, in the original Self-Torturer case, the cumulative result of his actions threatens to be self-defeating. That the cumulative effects of the self-interest action would be against the agent's self-interest suggests irrationality. By allowing the agent to turn back the dial for a sufficiently high price, we can turn the agent into a money pump. In the Beneficent Self-Torturer and in Distant Rescue, there is no similar suggestion of self-defeat or irrationality. There is, for example, no sense in which the cumulative effect of my individual donations puts me in a place such that I am now morally required to try to take back all that I have donated. It is doubtful that I would even be morally permitted to take anything back.

On account of this remaining difference between the Self-Torturer and the Beneficent Self-Torturer, I find it less convincing that the cumulative cost to me of my donations reveals that there must be something defective about the judgment that each of the individual donations is morally required. That is, the mere fact that I would not be required to turn the dial from zero to 20 to save 20 people does not show that I would not be required to perform each of a series of actions that cumulatively have the equivalent consequences both to myself and to the twenty. While there may be something surprising about this (particularly since there don't seem to be any reasons that favor advancing the dial incrementally rather than all at once), it doesn't, at least in any obvious way, involve any self-defeat or irrationality.

Despite these differences, Quinn's Self-Torturer may instruct us in two ways. First, the Self-Torturer suggests that, at least in the pursuit of our own self-interest, in our

assessments about whether we should perform some particular action (to advance the dial by one setting) we seem to have a blind spot for facts about our past and potentially future actions, and for how such facts can even be relevant. In the case of the Self-Torturer, such facts clearly are relevant, but we continue to have a difficult time incorporating their significance into any particular decision. This supports the suspicion that, in cases like Beneficent Self-Torturer and Distant Rescue, we may be affected by the same kind of blind spot. So, even if facts about our past and potentially future financial contributions are directly relevant to the question of whether we should or must donate another $10, we might expect that the significance of these facts will be difficult for us to see and appreciate when posing the question this way. Although we shouldn't conclude that there is anything misguided about the challenge to find counteracting considerations to keep the rescue principle from generating shockingly extreme demands in cases like Distant Rescue, Quinn's Self-Torturer provides us with a reason for suspicion: we may have difficulty appreciating the counteracting force of our past and potentially future contributions.

The second way in which Quinn's Self-Torturer is instructive is this: even though in Beneficent Self-Torturer the series of individual advances of the dial are not self-defeating, it is nevertheless quite plausible that the cumulative effect on the benefactor is so terrible that it couldn't possibly be a sacrifice that morality effectively demands, even if this demand arises via a series of smaller requirements (the individually trivial sacrifices each of which rescues an imperiled person). If, though, we so assess the Beneficent Self-Torturer, then however we should think about the counteracting consideration in that case may give us insight into how we should think about Distant

Rescue. In the Beneficent Self-Torturer, the salient fact obtains. If we can more confidently suppose that considerations of some sort prevent the rescue principle from demanding self-inflicted continuous electrical torture, then we can more confidently suppose two things: first that there is, at least in this case, a counteracting consideration for us to identify; and, second, that the counteracting consideration presumably has something to do with the way that costs of individual sacrifices accumulate for the beneficent agent (possibly in combination with other factors). At least at some point, then, the cumulative cost of individual responses to the salient fact can become so great that this can limit the would-be benefactor's obligation to respond to the salient fact (despite the individual cost of rescuing another individual person remaining trivial).

Unless more can be said, however, these lessons would only discredit the assumption that the cumulative costs to me of acting, or having acted, similarly do not factor into the question of how I must respond to an individual opportunity to rescue another person at a small cost to myself. That assumption, however, has already been dismissed. We rejected it in rejecting what Cullity calls the iterative approach (under its first interpretation).

Moreover, the remaining differences between the Beneficent Self-Torturer and Distant Rescue mean that it still remains unclear whether, in the latter case, the cumulative costs of responding to the salient fact can ever similarly limit what I must do in any individual decision about whether to rescue another person for $10. At least when the costs to me are primarily financial, there is reason to think that, at least in some cases,[34] most of what I posses can be required of me even to rescue a single stranger. If

[34] Peter Unger's "Bob's Bugatti" (see p. 134 below) is a vivid example of a high-cost rescue scenario, where the person to be rescued is a stranger to the potential rescuee (1996: 135-36). It may, of course, be

such cases show that even most of my personal fortune may have to be sacrificed in order to save one stranger, then it is less plausible that the accumulation of small dollar amounts sets a significant limit on my responses to the salient fact.

---

argued that we should separate off Bob's Bugatti as a special sort of case, or that the kind of sacrifice required in Bob's Bugatti is qualitatively different in important ways from the kind of standing obligation that threatens to arise in Heroes for Hamiltons (I pursue this line of thought in chapter 4, section 3.3). Other high-cost rescue scenarios are described in Rachels 1975: 79, Murphy 2000: 127, Hershenov 2002: 277, Ashford 2003: 274.

CHAPTER FOUR:
THE POND QUESTION AND THE FAMINE QUESTION

In chapter three I presented the differentiation question and I explained why the moderate is under pressure to find a plausible answer to it. In this chapter I consider what it is that makes such an answer so difficult to provide. Though proposed answers to the question are many and diverse, they all tend to generate dissatisfaction of one kind or another. Clarifying the difficulties provides guidance, both for constructing a more promising answer to the differentiation question and for acknowledging the kinds of dissatisfaction that are, for the moderate, inevitable.

Given a standing opportunity to rescue others in Distant Rescue, the moderate is committed to this: I must rescue the imperiled person in Shallow Pond even when I am not required to perform a rescue through Heroes for Hamiltons (the aid organization of Distant Rescue). The differentiation question asks why. As such, the differentiation question is a conjunction of two component questions, which I will call the pond question and the famine question:

> *Pond question*: Why must I rescue the person who would perish for want of my action in Shallow Pond?
>
> *Famine question*: Why don't I have to rescue anyone else through Heroes for Hamiltons?

The differentiation question's difficulties can be seen in the way that answers to each of the component questions, which appear fitting (at least apart from a certain angle), leave us in a poor position to answer the composite question. In this chapter I will try to bear out this idea. First I consider the seemingly fitting answers to the pond question and the famine question, and how such answers increase the difficulty of the differentiation question. I then discuss how an aim to mitigate these tensions provides the moderate with guidance in answering to the differentiation question. That guidance will allow us to propose the general shape of a promising moderate solution and to articulate the consequent challenges.

1. The Pond Question

1.1 Sufficient answers and inappropriate additions

When considered on its own (and by this I mean specifically apart from the differentiation question), a seemingly fitting answer to the pond question would be something like this: I must rescue this person because he/she is perishing, and I can so easily rescue him/her. That is, a seemingly appropriate answer will point to something close to the salient fact.[35] Put another way, we can say that, in the context of the Shallow

[35] It might be thought that there is a crucial difference between the salient fact and the suggested answer to the pond question: in the answer to the pond question, the demonstrative "this" picks out a specific person, whereas the salient fact obtains so long as *someone or other* gets rescued if I donate $10 to Heroes for Hamiltons—and, if I refuse to perform the rescue, no particular person perishes who would otherwise be saved. Caspar Hare provides a helpful discussion of this distinction in "Obligations to Merely Statistical People" (forthcoming), tentatively concluding that if the distinction does matter, it doesn't matter very much. For what it is worth, my own intuitions are that my obligations do not change when we modify either Shallow Pond or the operations of Heroes for Hamiltons so as to eliminate this difference. (Imagine, for example, that there are many imperiled people in the pond, and I only have the opportunity to rescue someone or other—and who it is who *would be* rescued by my action isn't determined. Alternatively, we can imagine that, after I open a tab on the donation page of Heroes for Hamilton's website, the website

Pond, and of pond-like cases generally, one need look no further than the salient fact to discover that one is obligated to perform the rescue. Plausibly this is because under such conditions the salient fact is a sufficient condition for the duty. That is, given the circumstances that constitute a pond-like case (whatever they are) the fact that I can rescue the imperiled person at a small personal cost is sufficient to guarantee my duty to do so.

This apparent sufficiency of the salient fact is further evinced in attempts to answer the pond question that cite the both the salient fact and certain additional features. If I were to say that the reason I am obligated to rescue the child in Shallow Pond is that I can easily rescue her *and* she is imperiled through no fault of her own, then I seem to say too much. Adding such additional information is not just unnecessary; it is inappropriate insofar as it suggests that I would not be so obligated if the child were drowning in the pond because of her own recklessness. In the pond-like cases there is a sense in which it is inappropriate for me even consider whether the person is imperiled through some fault of her own. The same is true for a wide range of further considerations I could conceivably contemplate in the pond-like case. Facing such an emergency, it is

---

generates a number, specifically for me, corresponding to a particular rescue candidate. If I close this tab before donating, that number gets shuffled to the bottom, nearly guaranteeing that the candidate won't be rescued; if, on the other hand, I donate another ten dollars, the particular candidate will be rescued and I am promptly issued the number of a new rescue candidate.)

Hare is inclined to grant some significance to the distinction inasmuch as he takes consolation in the fact that there is no one dead or seriously imperiled now who would have been rescued had he donated more money to charity. He reports this consolation in the face of the simultaneous (supposed) fact that some distant person would (likely) have been rescued if he had donated more money (13-14). I wonder if Hare's consolation—or at least the plausibility of its legitimacy—isn't in part due to an understandable fallacy: when I say to myself, "There isn't anyone whom my donation would have benefitted," it easy to take this to mean that my donation would not have benefitted anyone (i.e., my donation would not have made any positive and significant difference). Contrary to Hare, moreover, it is not as if refusing to give more money "was just very, very slightly bad for each of a vast multitude of people" (14). That Hare makes such a claim suggests that, in taking consolation, he temporarily loses sight of a more salient detail: that, by his own assumptions, giving more money would, in addition to being very slightly good for each of a vast multitude, likely have also been tremendously good for someone or other, saving such a person from sickness or death.

inappropriate to first consider why the child is imperiled in the pond, or whether or not rescuing such children might damage the socially important incentive for parents to be vigilant over the safety of their young ones, or whether or not I might be able to make up for letting this child drown by boosting my contributions to Oxfam, or to the Nature Conservancy.[36] This suggests that many such further considerations, some of which provide tempting answers to the famine question, do not properly figure into the explanation of why I must rescue the imperiled person in the pond-like cases.

There are, then, two features here. First, it appears to be entirely appropriate and sufficient to answer the pond question with the salient fact. Second, additional considerations appended to the salient fact in answering the pond question are often inappropriate. What is it that accounts for these features of the pond question? A tempting answer is that the salient fact simply is sufficient to require me to act: since the salient fact requires me to act wherever it obtains, it explains why I must act in any particular case in which it obtains. This is close to how James Fishkin understands the demandingness problems associated with the famine-relief argument. According to Fishkin, we normally assume that "if [one] can save a human life at minor cost, [one is] obligated to do so" (1982: 3, 65), with the problem, of course, that the salient fact also obtains in the famine-like cases, whose occurrence, or possibility of occurrence, may also be "normal."

[36] Douglas Odegard and Paul Gomberg draw attention to these observations (Odegard 1989: 81-2, Gomberg 2002). Noting the asymmetry with the sorts of considerations we take to be relevant to famine-like cases, Gomberg offers this as evidence that it is a mistake to suppose commonality in our reasons to respond to emergency cases like Shallow Pond and our reasons to respond to distant poverty (2002: 37-40). Nigel Dower suggests that Peter Singer's own argument, and perhaps utilitarianism generally, provides "the wrong kind of reason" for rescue, "by making this care impersonal because it is simply a contingent means towards maximizing the good, or…of preventing evils…." (2003: 647).

1.2 Avoiding Fishkin's inference

If Fishkin is right, and the assumption he ascribes to us is correct, then the demanding conclusion cannot be plausibly avoided. If Fishkin is right, but the assumption he ascribes to us is incorrect, then the moderate should not only explain why the imputed assumption is incorrect but also explain the widespread error. I see no reason, however, why the moderate should think that Fishkin is right. What Fishkin interprets as a general assumption may rather be this more restricted one: *for pond-like cases* we generally assume that if I can save a life at minor cost, then I am obligated to do so. Or, as I put it before: in the context of pond-like cases, the salient fact is, and is generally assumed to be, sufficient to show that I have a duty to perform the rescue. The moderate needn't grapple with Fishkin's more general assumption if it can just be explained why the salient fact seems to be a fitting and sufficient answer to the pond question, and why it is that it often seems inappropriate to append further considerations to such an answer.

One alternative explanation, tempting to the moderate, is to appeal to the urgency of the pond-like cases, or to say that in such cases the factors that might counteract the presumptive duty to perform the rescue are so unlikely that it is inappropriate even to look for them when considering what one should do. Consider first the urgency factor. It might be thought that the inappropriateness of adding further considerations to the salient fact is due to the special urgency often associated with pond-like cases.[37] The Shallow

[37] On the association with urgency, Murphy's remark is typical: "What we can call rescue situations are, paradigmatically, sudden, unexpected, serious harm-threatening emergencies that arise in the agent's immediate vicinity: the standard example is that of being easily able to save a child from drowning in a shallow pond" (2000: 127). A.I. Cohen (2005) proposes the feature of emergency is "a key, morally relevant difference between babies drowning at our feet and distant peoples suffering from hunger."

Pond is an emergency that calls for immediate action. Since there is no time to weigh the various particulars of the situation, considering various factors beyond the salient fact is inappropriate.

This urgency-based explanation of the fitting answers to the pond question loses its appeal, however, when we consider pond-like cases prospectively and retrospectively. We can regard Shallow Pond as a situation I might someday have to face, and therefore as a type of scenario for which I can develop a plan for what to do should the need arise. In the relaxed position of developing a plan, I readily concede that it would be good for me to at least be disposed to respond to such emergencies automatically, responding simply to the salient fact. Erring on the safe side, I will be better able to fulfill any rescue duty that I do encounter at the tolerable risk of possibly performing a small-sacrifice rescue even when I am not morally required to do so. This, however, is a dubious explanation of why I should be disposed to respond automatically to the pond-like case emergencies. It allows that in some pond-like cases I am not in fact required to perform the rescue.

Looking at the Shallow Pond retrospectively this point becomes still clearer. If I have earlier refused to rescue the drowning child, I do not think that, upon further investigation into the sorts of considerations listed earlier, it might turn out that my omission was permissible. If it turns out, for example, that the child was in the shallow pond because of personal irresponsibility, I do not then judge that this made my omission permissible. Such facts would not imply that I was morally lucky.

This urgency-based explanation also overlooks the counterfactual possibility that my own deliberative capacities enable me to cultivate highly discriminating dispositions

concerning emergency rescue scenarios, so that, even if most people should settle with the ham-fisted disposition to automatically respond to all such emergencies, I am permitted to take a different approach, enabling me to both refuse to perform some pond-like case rescues and rebut any criticism that I have acted wrongly, deliberated poorly, or that I possess any faulty dispositions.[38]

The related explanation, that of saying that, in a world like ours, the presence of counteracting factors would be so unlikely in pond-like cases as to be appropriately disregarded, can first be illustrated. In Shallow Pond, the presumptive duty to perform the rescue could plausibly be counteracted if it were also true that, as an unintended consequence, my action would serve to perpetuate the societal conditions that are to blame for numerous yearly fatalities. Since, however, in the actual world it is so unlikely that pond-like case rescues have this side-effect, this is the kind of consideration which, when faced with a pond-like case, it is inappropriate for me to seriously consider. Going down the list of considerations that the moderate is prone to entertaining when considering the corresponding famine question, we might similarly preclude each as inappropriate for pond-like cases in a world like ours, leaving me with a compelling reason to respond simply to the salient fact in a case like Shallow Pond.

---

[38] The urgency explanation is further challenged by non-emergency variants of Shallow Pond. I might encounter someone in a remote wilderness who, though not currently experiencing intense, debilitating suffering, is going to do so next week, after which she will die. I can easily intervene to save her, but I have at least a week to decide (whether I intervene immediately or later makes no difference to her).

Finally, it is difficult to understand why we should not regard Distant Rescue (or at least close variants of that case) as also involving the element of urgency. The idea might be that, in Distant Rescue, a failure to send a donation just now isn't urgent because a failure to send the donation right now will simply mean that a different imperiled person gets rescued if you send the donation later. Since there is no reason to rescue one person rather than the other, there is no urgency to send the donation now. If, however, the assured later opportunity of someone else's rescue can defeat the apparent urgency of the opportunity of rescue now, then it is unclear why the availability of Heroes for Hamiltons wouldn't also defeat the apparent urgency in Shallow Pond (when combined with Distant Rescue): even if I don't act in time to save the drowning person right in front of me, I can always donate more later to Heroes for Hamiltons.

On this type of explanation, the moderate is claiming that the pond-like cases differ from the famine-like cases in an important way: the pond-like cases, but not the famine-like ones, constitute the sorts of circumstances for which it is highly unlikely that there are any factors that counteract the presumptive duty to perform the rescue. This difference could explain why we are inclined to answer the pond question as we are, considering it inappropriate to seriously regard the sorts of considerations that the moderate might suddenly take seriously in any famine-like case.

This route seems promising for the way that it also suggests an answer to the famine question. Our typical opportunities to support humanitarian efforts occasion strong doubts as to whether the salient fact even obtains. As often as not, we learn that the relevant aid organizations must rely on unreliable intermediaries, and it is typically not true that my small individual donation will have any such significant effect on what a large and efficient international aid organization accomplishes. Or, even if the salient fact isn't in question, there are still legitimate concerns about the deleterious side-effects of international aid efforts generally, or of my donations to any one organization in particular. We aren't surprised to learn that well-intentioned disaster relief efforts end up serving corrupt or unsustainable local practices. When I open my mailbox, it is difficult to resist the suspicion that my small donation just supports future advertisements, possibly drawing funds away from organizations seeking to provide long-term solutions, or to save the Atlantic bluefin tuna.

So, in answering why I needn't perform the rescue in Distant Rescue, the moderate would hope to identify some counteracting consideration in that case whose occurrence is also highly unlikely in the pond-like cases, thereby answering the

differentiation question. This very hope, however, is threatened by Distant Rescue's idealizations. The idealized famine-like case with which our moderate must deal is constructed so as to preclude many of the sorts of factors that might plausibly block or reduce one's obligations in the messier sorts of international aid opportunities usually encountered. All such plausible suspicions about the salient fact or the presence of counteracting considerations are eliminated in Distant Rescue. So, the suggested explanation of our answers to the pond question suggests a moderate strategy that is ultimately problematic. Viewed generally, the problem is that the suggested solution relies heavily on facts about how the world is, or usually is, and these are the sorts of facts that can't be counted on in the idealized cases that the moderate must ultimately confront in answering the differentiation challenge.[39]

We began with the pond question, noticing that a simple appeal to the salient fact is a fitting and ordinarily sufficient answer. Moreover, we noticed that answers that append further considerations to the salient fact often seem inappropriate. The challenge for the moderate is to explain these observations without resorting to Fishkin's idea that we problematically assume that the salient fact just is a sufficient condition for generating a duty to perform the rescue. The last explanation we considered gained plausibility through its promise to also provide a compatible answer to the famine question, and therefore, the differentiation question. The problem, however, is that its suggested

[39] This is of course not problem for someone who accepts the demanding conclusion in highly idealized cases like Distant Rescue, only resisting the demanding conclusion in the arguably distinctive famine-like cases actually encountered. Such a person may conclude that the moderate commits a fallacy when extending to the highly idealized cases her resistance to the demanding conclusion—resistance that is appropriate for the more familiar, actual cases. On this view, the moderate commits a fallacy of over-generalization. Though this would not be a moderate position in my defined sense, it strikes me as a relatively attractive alternative, despite the concerns mentioned in chapter 3, section 2.4. The attractiveness of this alternative, however, somewhat depends on the success or failure of my attempt to justify the more thoroughgoing rejection of the demanding conclusion.

answers are endemically threatened by the idealized cases with which our moderate must contend.

Since the difficulty here concerns the suggested answer the famine question, I will now consider that question directly.

2. The Famine Question

The famine question asks why I am not similarly required to rescue a person through Heroes for Hamilton. The answer here, I think, will need to be sensitive to the force of the salient fact. The other factor that is likely to figure in centrally to the moderate's answer is the potentially high cumulative cost of compliance, to the benefactor, of any requirement to so respond to the salient fact. Call this factor *the costs of repeated rescue*. The problem is that this factor fails to distinguish the pond-like cases from the famine-like cases. So, while pointing to the costs of repeated rescue may seem to be the moderate's most straightforward explanation for why one isn't morally required to respond to the salient fact in Distant Rescue, the costs of repeated rescue fail to explain why we should treat Shallow Pond any differently. It fails to explain why, in the face of Distant Rescue, the imperiled person in Shallow Pond isn't to be regarded as just one other person who can be rescued at a very small cost to oneself.

If, however, the costs of repeated rescue are most plausibly motivating the moderate to distinguish between Shallow Pond and Distant Rescue, then such costs will likely play a role in the most promising justifications for the moderate's position. Attempts to defend the moderate's position without appealing to the factor actually

motivating that position will tend to generate inadequate rationalizations of only some of the moderate's judgments rather than justifications of them all.

2.1 The centrality of the costs of repeated rescue

To defend these conjectures I need to argue that the costs of repeated rescue are an important motivating factor in resisting the demanding conclusion. In the last chapter I tried to spell out a sense in which the moderate is pressured by his or her own will to find a solution to the differentiation challenge. These pressures arise in part because of the severe costs of repeated rescue arising if Distant Rescue obligates one to perform a rescue whenever Shallow Pond does. There, my autobiographical descriptions of those pressures were meant to ground a measure of skepticism towards the moderate's position. The high costs of repeated rescue plausibly explain, without in any obvious way justifying, the moderate's insistence that Distant Rescue is relevantly different from Shallow Pond. This ready availability of a potentially debunking explanation undermines the moderate's optimism that there is, or must be, a yet-to-be-identified solution to the differentiation challenge.

That potentially debunking explanation of the moderate's position was supposed to show that, even if the moderate lacks other legitimate grounds for rejecting the demanding conclusion, the consideration of the costs of repeated rescue would continue to motivate that rejection. My present aim is different. Here I want to argue that insofar as the moderate's position has promise, the costs of repeated rescue are in fact motivating that position. I start by defending a modest claim about motivation:

Weak motivational claim: our[40] resistance to the demanding conclusion is typically motivated by the perceived threat of the high costs of repeated rescue.

2.2 The student's progression of thought

To support the weak motivational claim, we can begin by noting what is plausibly a common progression of thought when a student is first confronted with Singer's famine-relief argument. It is common, I take it, for students to provisionally concede that we have a duty to rescue another person when this can be done at a very small personal cost. It is also common for students to concede that factors such as mere distance, or the number of other people involved, do not alter this duty. What frequently triggers a change of mind is the realization that, in our present world, such a duty threatens to make shockingly great demands of us given the accumulating costs of having to perform many such rescues. Once their minds are so changed, they are confronted with the task of explaining why one is not required to perform the rescues in the famine-like cases.

This common progression of thought provides a sense in which perceived costs of repeated rescue often motivate the moderate's efforts to answer the differentiation question. This anecdotal story gains plausibility on five counts. First, such a story might fit with one's own self-examination into why one seeks a moderate position, or is at least tempted toward one. Second, it fits well with how others have described their resistance

[40] A note about my use of the first-person plural. I am trying to identify a plausible defense of the moderate position. Since I sympathize with that position, and my project is to understand these sympathies, I find it helpful to argue from the first-person. As a condition on plausibility, I take it that the defense cannot rely on my own highly idiosyncratic judgments. I therefore try to restrict my appeals to my own intuitive judgments to those which I expect will have broad appeal, and which will be shared by many others, and especially by those attracted to the moderate position. For this reason I use the first-person *plural*. If it turns out that these expectations are credibly disputed I might revise my claims to the first-person singular, and propose more modestly that I am seeking a moderate position that is defensible given the judgments that *I*, as someone who has given a lot of thought to these issues, find highly plausible and difficult to dismiss.

to the demanding conclusion.[41] Third, given a person's natural aversion to unexpected and surprisingly burdensome personal responsibilities, the threat of otherwise enormous costs of repeated rescue provokes the right kind of motivation. Fourth, it would not be mysterious if the cumulative costs to an individual of compliance with a supposed moral principle is relevant to determining the legitimacy of the principle (compare, e.g., the mysterious proposal that the legitimacy of a moral principle is sensitive to the question of how compliance would affect the lunar landscape). Fifth, there is the difficulty of finding any other promising answers to what is driving our resistance. In resisting the famine-relief argument and its implications, one might propose any number of factors other than the costs of repeated rescue to explain or to justify that resistance. It is quite common, however, to discover that one's resistance to the famine-relief argument and its implications outlives any commitment one has to such proposals.

2.3 A normative lens

The weak motivational claim doesn't claim quite enough about the centrality of the threat of high cumulative costs of compliance in motivating resistance to the demanding conclusion. I would like to suggest further that, if it were not for those high cumulative costs, the moderate would likely not feel the need to differentiate Distant Rescue from Shallow Pond in terms of moral obligations. Consider, therefore, a stronger claim:

> *Strong motivational claim*: we would typically not be motivated to resist the demanding conclusion apart from the threat of the high costs of repeated rescue.

[41] For example, Herman 2000: 42.

While the weak motivational claim lends plausibility to the strong motivational claim, we cannot infer the stronger claim from the weaker. The stronger claim requires something more than the plausible story about the student's progression of thoughts. Even if true, the story about our progression of thought doesn't of itself suggest that we wouldn't end up with similar judgments about the normative dissimilarity of Distant Rescue and Shallow Pond independently of the consideration of the costs of repeated rescue. Such observations about our actual progression of thought are consistent with the view that additional factors would motivate us toward a moderate's position if consideration of cumulative costs didn't kick in. If, for example, we were blinded to the costs of repeated rescue, we might nevertheless be pressured by other factors to adopt the moderate's position.

So it is difficult to see how we might demonstrate the truth of the strong motivational claim. We would need to rule out all factors other than the high costs of repeated rescue that might independently trigger our resistance to the demanding conclusion. I therefore look for a way of narrowing the field of candidates. I do this in response to what the range of these options might otherwise be, the inevitable concern that I might have failed to imagine them all, and the tedium of individually ruling out all the factors that I do succeed in imagining.

Given my larger project, not all the factors that might falsify the strong motivational claim interest me. My larger project is to locate the most plausible defense of the moderate's position. The moderate's position resists the demanding conclusion in the famine-like cases. Since I am looking for a plausible *defense* of that resistance, and not just a plausible explanation of it, only some of the factors that might motivate our

resistance to the demanding conclusion are of interest. Since, moreover, any factor that could justify resistance to the demanding conclusion would also motivate that resistance (given the pressure to resist the demanding conclusion, explored in the last chapter), I am only interested in a subset of the factors that might render the strong motivational claim false. I am interested in those factors that could show a different claim to be false:

> *Justification claim*: we cannot plausibly justify our resistance to the demanding conclusion independently of the high costs of repeated rescue.

The move from the strong motivational claim to the justification claim reminds us that, of all the factors that might explain our resistance to the demanding conclusion, we are primarily interested in those factors that might also justify our resistance. The factors that actually motivate our resistance should figure in to the justification of that resistance if that justification is to both accurately track our resistance and avoid the appearance of a mere rationalization.

With this last move we bring a normative lens to what up until now has been a descriptive project. This lens makes some of the factors that might explain our resistance to the demanding conclusion more worthy of investigation than others.

2.4 The pond asymmetry

The next step might be to begin a process of elimination, considering whatever factor—other than the threat of high costs of repeated rescue—comes to mind as a promising reason for resisting the demanding conclusion. Appeals to analogy are a popular method for eliminating factors. While many proposed factors seem to be defensible reasons for resisting the demanding conclusion in Distant Rescue, such factors intuitively fail to license any similar permissions when added to pond-like cases. Think,

for example, of the fact that, in Distant Rescue, there are other people who are better positioned to perform the same rescues, or that the primary responsibility for dealing with such needs falls to other people or institutions. If we build these features in to the pond-like case, they intuitively won't provide permissions for refraining from rescuing.

Given the great variety of differences between Distant Rescue and the pond-like cases, the mere fact that a putatively significant feature in Distant Rescue seemingly fails to carry the same sort of significance in a pond-like case cannot, for reasons previously discussed, demonstrate the feature's insignificance. The feature's apparently changing significance might arise from its interaction with other variables that differentiate Distant Rescue and Shallow Pond. This warns against hastily eliminating factors through appeals to analogy.

With this weakness of the appeals to analogy in mind, alongside the inconvenience of still having too many potential factors to address, I introduce the *pond asymmetry*: there is a set of special features of Distant Rescue which, when added to a pond-like case, seem to license the same kinds of permissions that the moderate finds appropriate in Distant Rescue. Consider an example.

> *Nightmare:* The scenario is like Shallow Pond except that, instead of one child drowning in the pond, there are multitudes of people drowning, and there are multitudes of such ponds. Some volunteers are here to help, but they are far too few. After helping a few people get to shore, I see that if I were to rescue all of the people who are drowning in this one pond, I would at that time have reached the back of the pond and I would then be standing directly in front of another pond. Wiping the sweat from my brow, I observe that a series of such ponds stretches into the horizon. I recall having heard about this depressing state of affairs on NPR, but somehow it never stuck with me. Now, for the first time, I am witnessing the horror firsthand, and I am amazed that people don't find it more action-worthy. Confirming the reports, I see the canals transporting hundreds of dead corpses towards the sea while small tornados continuously deposit new victims, having raptured them up from afar. Each new person, gently dropped from the air, dizzily flails about in the shallows and then drowns.

My name for this case is a nod to Garrett Cullity's "nightmarish scenario." Here is what Cullity says about a scenario of just this kind.

> It is surely far from intuitively obvious that it would be wrong of me not to spend practically every waking moment saving lives….No doubt, saving no one would be wrong, but would it be obviously wrong, say, to spend my mornings pulling people out of the water and my afternoons pursuing my own life?" (1996: 61)[42]

The pond asymmetry enables us to divide into two groups the factors that might explain our resistance to the demanding conclusion in Distant Rescue. In the first group are those factors, or set of factors, that, though they could explain our resistance in Distant Rescue, fail to prompt any similar resistance when they are added to a pond-like case. In the second group are those factors, or set of factors, that not only explain our resistance in Distant Rescue, they also prompt a similar resistance when they are added to a pond-like case. Given the greater explanatory power of the factors in the second group, we now have further reason to dismiss the factors of the first group as plausible explanations of our resistance to the demanding conclusion in Distant Rescue. As such, the pond asymmetry gives us a promising way of further limiting our search. If there is a factor, or set of factors, that explains that resistance, it is most likely going to be found among the features that contribute to transforming the Shallow Pond into Nightmare.[43]

---

[42] Neera Badhwar shares Cullity's response, attributing a morally significant difference between the pond-like cases and the famine-like cases on the chronic nature of the needs in the latter: "If emergencies were truly like chronic evils in being an ongoing feature of life, who would not say that although drowning children were tragic, there was nothing wrong in not spending one's life saving them?" (2006: 81-2). She suggests that this judgment is reflected in "the actual attitudes of people in poor countries, such as India, toward the poor. These attitudes are shared both by people who, though not destitute themselves, live cheek-by-jowl with the destitute, and by people who being relatively well-off or rich live physically removed from the destitute but within daily sight of them" (83).

[43] I make only a claim of likelihood because it still possible that what explains our resistance to the demanding conclusion in Nightmare is entirely different from what explains our resistance to the demanding conclusion in Distant Rescue.

2.5 Applying the constraints

We are looking for factors that explain our resistance to the demanding conclusion. We have constrained the search to those factors that are distinct from the threat of shockingly high costs of repeated rescue, that might plausibly help to justify that resistance, and that are common to both Distant Rescue and Nightmare.

These constraints preclude many of the proposals found in the literature, so long as they are supposed to justify resistance to the demanding conclusion independently of the threat of the high costs of repeated rescue. Since we have constrained our search to those factors present in Distant Rescue, we can eliminate many other proposals for rejecting highly demanding obligations to those imperiled in famine-like cases through the built-in idealizations. We can dismiss as unpromising still more of the remaining factors because some of them are absent in Nightmare.

2.6 The alignment test

After so limiting the field of candidates that promisingly explains our resistance to the demanding conclusion, we are near to the point at which the remaining options must be directly considered. The reasons for dismissing these remaining options will vary from case to case. Before engaging in this task, I want to explain a generally useful evaluative tool.

For some of the candidate proposals I observe that, while they would offer reasons for resisting the demanding conclusion, they do so without quite providing reasons for *my* resistance to the demanding conclusion. Since I expect that other moderates will share many of the nuances in my resistance to the demanding conclusion,

I tentatively regard my resistance to the demanding conclusion as *our* resistance to the same.[44] Our resistance to the demanding conclusion has a content that goes beyond simple resistance to the claim that, in Distant Rescue, one is morally obligated to perform the rescues so long as the salient fact obtains. Our resistance is marked by insensitivity to a variety of small adjustments we might make to Distant Rescue. If we try to explain our resistance to the demanding conclusion in terms of a given factor, we can often test whether or not that factor accounts for these nuances in our resistance. If a proposed factor for explaining our resistance fails to account for these nuances, we can say that it fails the *alignment test*.

By qualifying these adjustments as *small*, I mean to contrast the alignment test to what I earlier called the "appeal to analogy." Appeals to analogy often involve extensive transformations of the analogized case, changes that would take a case like Distant Rescue and transform it into a case resembling Shallow Pond. Since such changes are far more extensive they introduce stronger grounds for suspecting that somewhere among those changes other normatively significant factors have been introduced, factors that might plausibly allow alternative accounts for the resulting intuitive moral judgments we have about such cases.

Merely adding one feature to Distant Rescue is sometimes enough to show that a given factor fails to adequately account for our resistance to the demanding conclusion: our resistance may persist even though the supposedly crucial difference is narrowly eliminated. If, for example, we add to Distant Rescue the stipulation that all of the people whom we might rescue are imperiled only because of an unforeseeable natural disaster, emergency preparations for which could not have been expected of any

[44] See footnote 40 about my use of the first-person plural.

institution, then our resistance to the demanding conclusion is likely to persist. This persistence casts doubt on attempts to explain our resistance to the demanding conclusion merely in terms of peculiar causes of the needs.[45]

We need to proceed with caution in considering proposals in this way. Any elimination of a given factor will be tentative for at least two reasons. First, the possibility remains that our resistance to the demanding conclusion, with all of its nuances, is not to be justified in its entirety. It may be best to accommodate some revisions to our initial intuitive judgments. Although an explanation that accounts for much of our resistance to the demanding conclusion might contain counterintuitive implications, it may be best to accept those implications. Second, even if our resistance is to be justified in its entirety, every aspect of that resistance needn't be explained in terms of the same factor(s). Our resistance to the demanding conclusion in Distant Rescue may be justified in terms of one factor while an entirely different factor justifies our similar resistance in the slightly revised scenario in which the first factor is absent. Although such a possibility tends to lose plausibility as we increase the similarity of two test cases, it is a possibility to bear in mind when applying the alignment test.

With these caveats, I consider some of the proposals that remain.

2.7 The possibility of more ideal collective responses

A prominent class of the remaining proposals seeks to justify resistance to the demanding conclusion in terms of a distribution of the responsibilities for meeting these needs. These conceptions of how responsibilities are to be distributed are, or can be,

[45] John Kekes (2002) and Dale Jamieson (2005) suggest that our resistance to the demanding conclusion might appeal to such root causes.

independent of the shockingly high costs of repeated rescue. I begin with Joel Feinberg's proposal, which also serves to introduce several prominent variants.

2.8 Feinberg on the relevance of the more ideal collective solution

Although Joel Feinberg's resistance to the demanding conclusion may have been originally motivated by the threat of shockingly high costs of repeated rescue, those costs are not in any obvious way essential to his justifications for that resistance. For Feinberg, what is central is the possibility of a more ideal response. Feinberg considers Lord Thomas Macaulay's discussion of an Englishman in India circa 1835:

> [A] resident Englishman would encounter hordes of beggars on every street corner, and thousands of these would die every year of malnutrition or starvation. It was manifestly absurd to hold that each time a wealthy man encountered a beggar, he had a duty to rescue him by making a small contribution. (1984: 66)

Feinberg is discussing the concern that any way of specifying the Englishman's legal[46] duty to aid others will either rely upon an arbitrary line dividing duties from non-duties, or it will legally require a person to perform acts of rescue well beyond anything that morality could plausibly require. However, in addressing Lord Macaulay's scenario, Feinberg observes,

> Any practical, fair-minded person…[who is] also humane…would desire instead some sort of scheme of coordination that would allow the starving as a class to be rescued by the wealthy as a class without unjust enrichments of the unworthy or unfair disproportions in the contributions exacted from the donors. A modern state's welfare system, with its maintenance of an income floor for indigents paid out of the taxes from those able to pay, is just such a system. (66)

[46] Although Feinberg's concern is often directed toward the potential legal distinctions, he regards such legal distinctions as themselves grounded in *moral* distinctions. The objections to the legal distinctions Feinberg is considering are objections directed at questions of moral legitimacy.

In contrast to this scenario, Feinberg notes that there are different sorts of rescue scenarios in which the random and unpredictable emergencies of life that require time and effort, rather than money, from chance passersby, are not obviated by state welfare systems (66). Feinberg believes that the distinction between these two types of cases provides a principled way of distinguishing those cases in which we have a strict duty to provide aid, and those cases in which we do not. Strikingly, Feinberg proposes that the distinction holds even when the state welfare system is not up and running. So, for Feinberg, the Englishman has no strict duty to aid the hungry beggar in front of him even when there is no welfare system. For Feinberg, the salient fact is that the beggar's hunger *would be* better handled by a welfare system.[47] Feinberg's thought, that facts about the more ideal solution are so morally significant, appears to be motivated by his idea about the role or function of morality. For Feinberg, morality (or at least "public morality") "is essentially a way of coordinating private efforts for common goals" (1984: 67, 1987: 170).

It is fairly easy see what Feinberg would say about Distant Rescue and Nightmare. In each scenario there are presumably more ideal ways for the societies to deal with the problems. In Distant Rescue, perhaps the more ideal solution might involve large-scale development projects, the creation of tax-supported international emergency

[47] Feinberg dramatically illustrates his position in terms of the responsibility to put out fires:

> The reason we have the duty to report the fire but not the duty to fight it is not just that there is minimal effort required in the one case and not in the other. It is rather that the very strict social duty of putting out fires is most effectively and equitably discharged if it is split up in advance through the sharing of burdens and the assigning of special tasks. (1984: 68)

Here, again, the reason why the individual is not obligated to fight the fire is not that the task *will be* handled by a capable fire department, or that there is some fire department which is responsible for handling the fire; it is rather that firefighting *would be* "most effectively and equitably discharged" *if it were* addressed in a more organized way.

funds, political reform, and the cessation of wars. In Nightmare, the ideal solution might involve the National Guard and a community-funded bulldozer fleet to drain the ponds.

There is something promising in Feinberg's observations. The fact that there is a far more ideal cooperative solution to a common problem plausibly should affect the moral responsibilities of individuals. Evaluating Feinberg's proposal, however, is difficult because he doesn't spell out in sufficient detail what he has in mind when he says that one solution would be more ideal than another. If, as it often appears, what Feinberg has in mind are solutions that avoid imposing highly costly responsibilities upon any one person, then he likely isn't offering a proposal that is independent of the threat of those high costs of repeated rescue.

2.9 Rawls's division of labor

Feinberg's proposal has affinities with Rawls's emphasis on a distinction between the responsibilities that fall to institutions and those that fall directly to individuals. Having described such a distinction as a "division of labor" (1993: 268), it suggests a way of limiting the responsibilities that apply to individuals.

Samuel Scheffler argues that where Rawls explicitly uses the division-of-labor metaphor, he is contrasting, on the one hand, rules that apply to the basic structure of society for the purpose of securing the background conditions in which individual transactions can take place fairly, and, on the other, the rules that are to be followed by individuals concerning their own transactions with each other. Rawls's concern here regards the way in which seemingly just transactions among individuals can lead over time to unjust conditions. Such outcomes cannot be feasibly avoided by insisting that

everyone comply with norms applying directly to their individual conduct. Such rules would be bound to "exceed the capacity to grasp and follow them with sufficient ease," be "too complex," "require too much information to be correctly applied," or "enjoin individuals to engage in bargaining with many widely scattered third parties" (1993: 268-69). Rawls's solution is to apply the principles for maintaining background justice to the basic structure.

Since this division of labor is motivated by the impossibility of securing the value in question by rules applying to individuals (Scheffler 2005: 238-40, 244-45), this creates a significant disanalogy between Rawls's division of labor and the sort of division of labor that might directly justify a reduction in individual responsibilities for the distant needy. If background justice isn't secured at the institutional level, then there is by hypothesis nothing that an individual can do to remedy the problem. When, on the other hand, the relevant institutions fail to secure distant strangers from peril, there is still much that the individual can do through Heroes for Hamiltons. Individuals are free of the responsibility to follow norms that might secure background justice because the relevant norms would be too complex to follow. No equivalent excuse is necessarily available concerning norms directing individuals to rescue the distant needy when the opportunities to do so are present and obvious.[48]

[48] Besides the contrast just mentioned, Scheffler identifies another distinction in Rawls between the rules applying to the basic structure and rules applying elsewhere. With this other contrast Rawls is emphasizing the limited scope of his project by acknowledging the diversity of values and principles that presumably would be relevant were his subject not restricted merely to the basic structure of society. While Rawls doesn't himself use the division-of-labor metaphor for this further distinction, Scheffler does (2005: 239). Here, however, the metaphor would most directly apply to the division of *theoretical* labor, whereby the task of identifying normative principles gets divided up into more manageable parts, justifying Rawls's focus on the basic structure of society without the need to comprehensively address the entire domain of justice and morality.

Used in the way that Scheffler and others intend,[49] however, the division-of-labor metaphor is rather meant to imply a way of portioning out responsibilities between institutions and individuals for securing just outcomes or realizing moral values, and, borrowing from Rawls's actual use of the metaphor, to do this in a way that leaves "individuals…free to act effectively in pursuit of their ends and without excessive constraints" (Rawls 1993: 268). Used in this way, the metaphor suggests that the responsibilities falling to an affluent individual for direct action on behalf of the distant needy are reduced on account of the fact that related responsibilities are apportioned to institutions. Rawls would ascribe some such responsibilities both to the basic structures of the societies to which the imperiled peoples belong, as well as to all the "well-ordered peoples" that owe a duty of assistance if these societies are "burdened" (1999: 106-113).

So employed, however, the division-of-labor metaphor is dubious. Extending a point once made by Robert Nozick, we shouldn't assume a model in which responsibility exists in a fixed amount to be rationed out. "Responsibility," he notes, "is not a bucket in which less remains when some is apportioned out" (1974: 130). To say that the responsibilities of the imperiled person's own society relieves me of the personal responsibility to directly rescue him sounds a lot like saying that the responsibilities of a child's parents relieve me from responsibilities of rescuing the child should she fall in the shallow pond (through parental negligence or otherwise).

[49] Thomas Nagel writes that "[t]he idea of a moral division of labor between individuals and institutions is…the form of a solution" for the problem of satisfying norms of impartiality in the face of motivational constraints, and "each person's reason to lead his own life" (1995: ch. 6). Murphy (1998) uses the metaphor similarly. Scheffler (2005) cites both. Approving what he takes to be Rawls's "'division of responsibility' between individuals and the larger society," Arthur Ripstein writes, "On Rawls's view (or at least my appropriation of it) we owe duties of support to just institutions, including institutions that aid those in need. But where just institutions are in place, particular individuals owe no duties of aid to others" (2000: 751, see also 756-58).

Since, then, we should not assume what the division-of-labor metaphor implies, we need an account of how the division of responsibilities affects individual duties towards the imperiled. Barbara Herman and Liam Murphy offer alternative proposals for just such an account, neither of which necessarily relies on the threat of high costs of repeated rescue. I will consider Herman's suggestions first.

2.10 Herman on inherited obligations

Like Feinberg, Herman suggests that strict duties of rescue might be largely isolated to those cases in which a more preferable cooperative strategy isn't feasible.[50] For Herman, however, an essential aspect of the more preferable strategy is that in it, people are fulfilling their "primary" obligations. Herman suggests that the demanding conclusion can be resisted by appealing to the distinction between primary obligations and inherited obligations.[51] Herman's basic idea goes something like this:

[50] This point is also echoed in Gewirth 1987.

[51] To be fair to Herman, this distinction between primary and inherited obligations is but one of her resources for articulating and justifying a moderate view of what beneficence requires. Before discussing the distinction between primary and inherited obligations, Herman (1984) endorses a distinctively Kantian basis for the duties of rescue and beneficence. That Kantian basis includes two important ideas: first, that duties of beneficence and rescue are grounded specifically in the value of rational agency, and second, that our others-regarding duties are mitigated by our obligations to cultivate our own agency. One way for Herman to resist the demanding conclusion is with the claim that compliance with such putative responsibilities would impermissibly undermine one's own agency (2001: 238-45). Apart from the faithfully Kantian framework upon which this type of argument relies, these ideas are of less interest to me here because of their implicit dependence on the costs of repeated rescue.

A separate resource is Herman's view about the "central function" of morality, and, in particular, about the way that morality needs to fit into "ordinary, everyday life," "securing routine action" "as a matter of course." At least two important claims emerge from this view of morality's role: first, that a "reasonable morality is…not something we are endlessly at war with (like a diet), nor a distant goal toward which we direct substantial amounts of our energy"; second, that in functioning to serve us in everyday life, moral guidance does "not easily extrapolate into extreme situations" (2000: 29-32). These views about morality's role suggest to Herman that while there should be a duty to rescue in the pond-like cases, this is not because of some general principle that can also be used to define our duties in the famine-like case. The duty to rescue in the pond-like cases follows instead from morality's role in meeting the "pressing needs that fall outside of or in the crevices of the institutional framework of [welfare] support" (2000: 43). With the pond-like case so handled by a kind of special duty that comports with morality's putative central

1) Need-at-a-distance typically arises only because more local people have defaulted on their primary obligations.

2) Therefore, the individual obligations we have towards persons in need-at-a-distance are typically inherited obligations.

3) The scope and content of inherited obligations are significantly limited.

4) Therefore, our obligations towards persons in need-at-a-distance are typically significantly limited.

Let us grant (1) and (2) in order to focus on (3). Herman begins with the general observation that inherited obligations are often not the same as their primary counterparts:

> One may take on only an aspect of another's more general obligation….it may not be appropriate in some circumstances for the holder of the derived obligation to act; in some cases, the derived obligation is only to forestall some bad effect that would occur were the original obligation to be unmet. Some cases might generate a derived obligation to help restore the primary obligation holder to full functioning. One can imagine cases in which the obligation (or what it can require) weakens as it descends. (2001: 250)

It is true that we cannot simply assume that the inherited duty is the same as the primary one. That said, it is not obvious that the relevant inherited obligation to rescue will differ in a way that makes them more *limited*. After all, defaulting on primary obligations often only makes a problem worse, causing the inheritors of the responsibility to inherit a larger problem. Moreover, whenever the inherited obligation does differ in scope or content from the primary obligation, we expect to be able to appreciate the reasons for why this is so. This certainly seems to be true in the examples of inherited obligations that Herman uses to illustrate this point. For these reasons, something more must be said before we conclude that our obligations to the distant needy are less demanding because they are inherited.

---

function, Herman has then blocked any implications this might have for the famine-like cases, our duties for which can then be constrained according to her ideas about beneficence discussed above.

2.11 Respecting the priority

Herman says that we must "adequately respect the priority of the local, primary obligation." She seems to take this to mean two distinct things. On the one hand, Herman takes the priority of primary obligation to mean that inherited obligations to aid others should not increase the likelihood that the primary obligation holders will continue to default (2001: 250). On the other hand, respecting primary obligations is also supposed to allow us to favor our own primary obligations, and in particular those responsibilities that might *seem* to be supplanted by inherited obligations (253). That is, respect for primary obligations might license us to "expend resources on higher-function needs close to us" instead of directing those same resources to meet more pressing needs arising because others have failed to meet *their* primary obligations.

Both of these ideas appear to follow from the more general suggestion that we have reason to try to ensure that obligations are met by their primary holders in particular. So, if I have resources to meet one of two needs, then, other things being equal, the suggestion seems to be that I should meet the one for which I have a primary responsibility. Herman suggests that I should still favor my primary obligation even if, other things being equal, the other need is in some respect more pressing (perhaps more stands to be lost if that other need isn't met).

But why should this be? I cannot think of a case in which what I should do plausibly depends *only* on whether or not my responsibility is primary, rather than inherited. Often something else of normative significance attaches to the fact that one of the responsibilities is primary. I can, however, imagine a fanciful case in which this is

not so, where the other normative factors don't obviously interfere. Suppose that two biologists are asked to closely monitor two tortoises inhabiting opposite ends of a remote island. On their way the biologists divide responsibilities: one is to monitor the western tortoise, and the other is to monitor the eastern one. One biologist, however, commits suicide before reaching the island, thereby defaulting on her obligation. Is the consideration of primary obligation still relevant in the remaining biologist's decision of which tortoise to monitor? I think not.

We need to hear an explanation of the priority of primary obligations, and why we are to favor the fulfillment of primary obligations over inherited ones. Although Herman does provide some of this explanation, what she says seems to make her overall proposal less interesting, since, as it turns out, the normative work is done by the sorts of commonplace considerations that are already frequently mentioned with regard to helping the needy. That is, Herman reminds us that when we try to meet needs arising from the failure of sound local institutions, there is the added danger that we will interfere with the development of better local institutions, and that our "direct palliative aid" will have hazardous side-effects (2001: 250). There is also the fact that those with the primary responsibility for the needs are often in many ways closer to those in need. This means that they will likely possess more insights into the problem. Moreover, and perhaps just because of commonplace expectations, "those to whom we are closest are often those we can most easily harm by our neglect" (247). Once we see that it is these sorts of considerations that are doing the normative work, it seems that what we ought to respect is simply these considerations directly—rather than the supposed priority of primary obligations per se.

2.12 Respecting the content

One final suggestion from Herman is that "[i]nherited obligations…respect the content of the primary obligations they take over" (2001: 251). Her suggestion is this: when the primary obligation is an obligation to have sound, local institutions, the content of the inherited obligation is similar: it is an obligation to support or develop sound, local institutions. Herman is attracted to this proposal because it avoids the implication that what gets inherited is simply the obligation to *directly alleviate* the problems resulting from the lack of sound, local institutions. Herman wants those inheriting the obligation to be able to direct their resources to the root problems, rather than (only) towards the "direct palliative aid" of those presently suffering because of the lack of sound, local institutions.

This isn't a promising way to secure the conclusion. First, it doesn't fit very well with her prior warning that we can't simply assume that the inherited duty is the same as the primary one. Moreover, when a person defaults on a primary obligation, this is likely to change the circumstances such that it would be foolish for the inheritors of the problem to try then to do what those with the primarily obligation should have done in the first place. The captain of an oil tanker is responsible for staying alert while navigating the vessel. Should the oil tanker default on this obligation and spill a million barrels of oil we may inherit very different obligations (and these inherited obligations, moreover, may be far more onerous). A government agency may have a primary responsibility for enforcing regulations on deep-water oil drilling. If defaulting on this obligation leads to an environmental disaster, it would be odd to insist in that moment that the inherited

obligation cannot be more urgent or more directly palliative. These cases suggest that even if the primary obligation is to create or support sound, local institutions serving the local people, this does not imply that the inherited obligation isn't, first and foremost, an obligation to alleviate directly the suffering that arises from the lack of such institutions. So, if resources can be redirected away from the direct palliative aid and to the development of sound, local institutions, this won't be simply because the primary obligation was an obligation to have sound local institutions.

2.13 Murphy on responsibility only for one's fair share

More so than the others, Liam Murphy makes it clear that, on his view, what motivates and justifies the resistance to the demanding conclusion is not the high costs of repeated rescue per se,[52] but rather that a requirement to make such sacrifices would

[52] Murphy's view here is in part driven by his skepticism towards the over-demandingness objection, which might have alternatively been called the "overly-costly objection." According to Murphy, if the high costs of repeated rescue are prompting us to reject strict obligations to the distant needy, this cannot simply be because such obligations would be too demanding. For one thing, it is unclear whether demandingness can be measured in a non-question-begging way that still manages to capture what is intuitively relevant about demands. Second, even if demandingness can be measured, it is difficult to see how we could appeal to demandingness as a way of plausibly limiting one's obligations to the distant needy. Third, if we accept an intuitive understanding of demandingness, we will have to acknowledge that morality can sometimes be extremely demanding. We will therefore still need to explain why demandingness limits moral responsibilities in some cases but not in others.

I address the first and second objection in "Better Collective Solutions." I will discuss issues related to Murphy's third objection later in this chapter. Regarding the first objection, the basic problem, as I see it, is that Murphy is unwilling to give our intuitive moral judgments any credence if they run afoul of utilitarianism; such judgments cannot be taken to support a valid objection that utilitarianism is overly-demanding because, in establishing what is to count as overly-demanding, such judgments beg the question against utilitarianism. My project, however, is to try to accommodate the very judgments that Murphy prefers to dismiss. As such, I am willing to interpret the over-demandingness objection in the "question-begging" way that Murphy resists. In other words, when our intuitive moral judgments indicate that utilitarianism is making demands that are too severe (particularly with regard to opportunities involving rescuing strangers at an individually low cost), my project is to see what can be said on behalf of those judgments.

illegitimately require you to do more than your fair share of a collective responsibility.[53, 54] Murphy, like Feinberg, attributes significance to facts about more ideal collective solutions. Like Herman, the significant aspect of the more ideal collective solution is, for Murphy, that others are fulfilling their moral responsibilities. Murphy differs from Herman in his proposal for why needs arising from the moral failures of others may have a more limited claim upon us. For Murphy, the issue is not whether meeting any particular need is your primary or inherited obligation. This question doesn't even matter. What matters is the degree to which the overall cost of improving the welfare of people generally is sustained by the (concurrent) moral failures of others. To the extent

[53] Murphy is developing a view proposed earlier by L. Jonathan Cohen: one's "duty is only to play a fair part in the performance of what is collectively obligatory. Burdens…ought to be fairly distributed, and ceteris paribus, no-one is morally required to take on more than his fair share of a burden because someone else defects…." (1981: 76). For Cohen there is a further reason for this limit: a code of ethics that passes obligations on to others after the original bearer of that obligation has defaulted would undermine the original bearer's motivation to fulfill his/her obligation. Cohen's reasoning strikes me as dubious. The motivation to fulfill my obligation is supposedly undermined by the thought that, if I don't fulfill my obligation, someone else will. However, the likelihood that someone else will fulfill my obligation after I have defaulted isn't obviously affected by whether or not my moral obligation would thereafter be transferred to someone else. Perhaps Cohen assumes that moral obligations are obligations for which social sanctions are appropriate, but that social sanctions are inappropriate in this case because, if they are effective, then they will somewhat undermine the motivation of the original bearer of the responsibility. This bears affinities to Feinberg's idea that morality is "essentially a way of coordinating private efforts for common goals." Feinberg's idea seems to be this: since morality coordinates our behavior (or it is at least supposed to) moral obligations should be those rules which, if followed, would result in an effective resolution of the problem with the costs of resolving it equitably shared (1984: 67).

[54] This type of claim can be generalized. Even if we concede the claim about the student's progression of thought, it may be that what makes the high costs of repeated rescue so objectionable is not the high costs per se (morality can be highly demanding after all) but something else peculiar to Distant Rescue (or the famine-like cases generally). Let *X* be any factor other than the threat of high costs of repeated rescue. Two claims can be made about X: (1) X is normatively relevant because of the threat of the high compliance costs; (2) the threat of the high compliance costs is normatively relevant because of X (with the implication that X might ground an independent objection to the demanding conclusion). Our constraints dismiss proposals that involve claims of the first, but not the second type. Often, however, claims of the second type lose plausibility from scenarios in which the given factor is eliminated while the threat of the high costs of repeated rescue is maintained, as in the case of the many factors eliminated by the pond asymmetry.

that the (concurrent) moral failures of others sustain these costs, any given person's moral responsibilities for meeting them will be diminished.[55]

Roughly stated, Murphy's proposal begins with the idea that, under conditions of full compliance from now on, each of us is required to "keep benefiting others until the point where further efforts would burden us as much as they would help others." The sacrifices (both active and passive) that each person would sustain under full compliance with this principle determine each person's *fair share* of the sacrifices required to fulfill the collective project of beneficence. Invoking, then, his premise that individuals "are not required to take on the entire responsibility of a collective project, but only their part in it," Murphy appeals to the idea of a fair share to limit what can be demanded of any individual under conditions of partial compliance. Stated roughly, the resulting standard, which he calls the "collective principle of beneficence," "requires agents to promote the well-being of others up to the level of sacrifice that would be optimal under full compliance" (2000: 7).

In assessing Murphy's proposal, we can set to the side its distinctively utilitarian elements. Murphy's most distinctive contribution is separable from his utilitarian commitments. I divide that contribution into two main claims:[56]

---

[55] One upshot of this is that Murphy's proposal easily avoids the obvious sorts of counterexamples to which related fair share proposals might be liable, as in, for example, Rachels 1979 (162-63). Murphy's proposal does not lead to the counterintuitive result that if there are two children drowning in the shallow pond, and there are two potential rescuers, then each potential rescuer is only responsible for rescuing one child, regardless of what the other person does.

[56] Regarding these two claims, the influence of Murphy's utilitarianism will show itself most in Murphy's defense of the first, and specifically in Murphy's claim that beneficence has an "agent-neutral goal." While the second claim only states the general restriction on the demands of beneficence, Murphy's full description of what beneficence demands (his "collective principle of beneficence") would re-introduce utilitarianism, with a requirement that the agent "must do as much good as possible" within the limit, and with a measurement of each person's fair share (which requires an account of the "compliance effects"—the effects to each person's wellbeing under conditions of full compliance to the utilitarian demand upon each person to do as much good as possible).

1. *The collectivist interpretation*: Beneficence is a project whose responsibility fundamentally falls to everyone *as a group*.

2. *The full-compliance standard* (for assessing individual responsibility): Individuals are only responsible for fulfilling their fair share of any collective project, where this fair share is assessed under the assumption of the full compliance of everyone from now on.[57]

Together, these two claims offer a distinctive proposal for why we should resist the demanding conclusion. Insofar as beneficence is a collective project, it is unlike certain special duties and deontological constraints whereby the idea of allocating a shared goal into smaller individual responsibilities would not be applicable. Moreover, because the project of beneficence invariably suffers from widespread non-compliance, we should resist the temptation to attribute to ourselves and to others moral responsibility for more than the fair share of the overly large task that has been left to be done.

I will focus on Murphy's full-compliance standard for assessing individual responsibility.[58] On Murphy's view, in assessing a moral principle governing the conduct of individuals who are collectively obligated to a given goal, we are to consider how everyone's full compliance with that principle (from now on) would impact a person's wellbeing—both through that person's own compliance and through the full compliance of others. Murphy's second claim amounts to a constraint: such principles are never to require an agent to act in a way that would diminish the wellbeing of herself or anyone else relative to what that person's wellbeing would be if everyone were to fully comply with the same principle from now on. A principle that violates this constraint would,

[57] What I am calling Murphy's full-compliance standard is an introductory version of what Murphy refers to as the "compliance condition." The final version of Murphy's compliance condition is cumbersome (2000: 80-82); the added complexities not only obscure the intuitive attractiveness of the original idea, I think that they also deviate from what is attractive about the original idea.

[58] I discuss the collectivist interpretation in "Better Collective Solutions," June 2011 draft, pp. 35-42.

according to Murphy, impose upon at least someone more than his/her fair share of the group's responsibility. This means that fair shares are not specified in terms of who is supposed to do what, but rather in terms of the amount of burden or sacrifice borne by individuals through actions made towards fulfilling of the collective goal. It also means that the fair shares of such burdens are assessed under the idealizing assumption that everyone will fully comply from now on. Since many hands make light work, Murphy's view places significant limits on the amount of help that beneficence requires of anyone, particularly whenever widespread and ongoing irresponsibility/immorality is responsible for significant losses in wellbeing.

Although a complete analysis of Murphy's compliance standard should examine the motivations for it and the ways in which it deviates from the intuitive idea "that it is objectionable to expect agents to take up the slack caused by the noncompliance of others" (1993: 278; 2000: 76-77, 89, 93, 127), here I simply want to look at the ways in which its implications run afoul of intuitive, and arguably common-sense, moral judgments. We can divide these implications into three categories.

Into the first category are the circumstances in which Murphy's compliance condition results in surprising permissibility. Whenever we have an agent who is already badly off, beneficence may require no actions from that person—even when that agent is the unique person who can help, even when that help is needed to meet the basic needs of the potential beneficiary, and even when the cost of helping would be extremely slight to the would-be benefactor. Murphy suggests that this badly off agent from whom beneficence never requires any sacrifice "describes the situation of many—perhaps most—people in our actual world" (2000: 87, see also 128, 131ff). As well as the badly

off, we can place in the same category any agents—including well-off ones—who have already fulfilled their fair share of the collective project of beneficence. Beneficence will likely not require that these people to do anything in the same scenarios (1993: 290 fn. 41). It should be noticed that this tendency of Murphy's view to demand too little does not require the potential benefactors to be badly off in any severe way, or to have already performed peculiarly great feats of beneficence. Such extremes are entirely unnecessary so long as either the number of non-compliers is sufficiently great or the neediness that is sustained by noncompliance/immorality is sufficiently high. War, or even war's threat, can easily create such circumstances (1997: 85; 2000: 124-26).

Into a second category, I place the scenarios in which Murphy's compliance condition doesn't appear to block the requirements of beneficence from becoming intuitively excessive. In this category are the scenarios in which there is a high proportion of impoverished people whose needs are not sustained by the ongoing noncompliance of others. So, for example, by simply increasing the inhabitants of poor nations and impoverishing them through natural disasters or *past* noncompliance, then the relatively fewer well-off citizens may still face demands of beneficence as severe as anything Peter Singer might be inclined to ascribe in his most radical moments. Murphy's compliance condition would be of no help in limiting demands in such cases (2000: 100-101).

In my third category belong those implications that give rise to what we might call, following Tim Mulgan, the "wrong facts objection" (Mulgan 1997; Murphy 1997: 83ff). The problem, already visible, is that whether or not an agent is required by beneficence to act seems to depend on the wrong sorts of considerations, and seems to be

insensitive to seemingly relevant ones. As we saw above, whether or not an individual is required to help another person depends upon how many other people are in need of help, and upon how many of these needs exist because of present or expected noncompliance. For the particular person needing help right in front of you, whether or not you are required to help him may depend on how many impoverished people are living on the Indian subcontinent: you may be required to help if they number 200 million, but not if they only number 20 million. You have also got to consider the number of present noncompliers. Or, even if the impoverished number 200 million, the existence of 50 million affluent noncompliers in Europe may again show that beneficence does not require you to help the needy person right in front of you (or, of course, anyone else).[59]

2.14 Conclusions

Doubtless there are other proposals, or other possible proposals, that both slip through the constraints I have outlined and attempt to justify resistance to the demanding conclusion in Distant Rescue independently of the repeated costs of rescue. My modest

[59] Torbjörn Tännsjö finds the idea of fare shares of a collective responsibility a promising way to explain our motivation to reject the demanding conclusion, finding evidence of this in a thought experiment, which adds a twist to Peter Unger's scenario of the Vintage Sedan (when you must decide whether to transport an injured person to the hospital in your car, knowing that this will ruin your expensive upholstery; if you don't so transport the person, he will certainly lose his leg):

> If I and five other persons find a person in distress, and if there is a simple way of helping him (such as standing together on a certain platform), provided we do all cooperate, and an expensive way of helping him (with a cost to the agent comparable to the cost to the Sedan owner, if he would have the upholstery of his car ruined), provided each act on his or her own, and if I do my share in the cooperative venture (I stand on the platform and urge the others to join me), then I think most people would say that I do what I ought to do. They say so even if my behaviour means that the person is not saved, because the other folks refuse to cooperate. Blame them, people would say, not me. (2000: 119)

Tännsjö's scenario would be more relevant to Distant Rescue if he knows with certainty that at least one of the other five will not cooperate. Is he still acting permissibly if he chooses to stand on the platform, knowing full well that this will mean that the injured person will lose his leg?

ambition is not to demonstrate the necessary failure of all such proposals, but rather to convey the lack of promise in this direction. Having pursued several of the prominent alternatives, I propose that the moderate should tentatively accept the justification claim, conceding, as a working hypothesis, that the intuitively fitting answer to the famine question is basically correct: the reason I am not required to provide another person-rescuing donation in Distant Rescue concerns the cumulative costs to me of repeatedly having to do so.

3. THE COMPOSITE QUESTION

The intuitively fitting answers to the differentiation question's two component questions leave us in a poor place for answering the composite question. I must perform the rescue in Shallow Pond simply because I can easily rescue an imperiled person. This doesn't, however, distinguish the opportunity presented in Shallow Pond from the opportunities presented in Distant Rescue. From the other direction, I am permitted to refuse rescuing the imperiled person in Distant Rescue because of the costs of repeatedly having to do so. Given Distant Rescue, however, the costs of repeated rescue fail to distinguish such rescue opportunities from the opportunity presented in Shallow Pond.

Beyond leaving the differentiation question unanswered, the difficulty of answering the differentiation question is increased by these intuitively fitting answers. In Shallow Pond, I am inclined, on the one hand, to regard the imperiled person simply as someone whom I might easily rescue. Adding further specifications to explain why I must perform often seems inappropriate. On the other hand, it is also clear that I cannot

regard the imperiled person as simply someone whom I might easily rescue, since this doesn't distinguish that person from the countless others who share that status in Distant Rescue.

Put another way, if the salient fact provides decisive reason to act in the pond-like cases, this gives some reason for expecting the salient fact to function in the same way in the famine-like cases. From the other direction, if the costs of repeated rescue counteract the salient fact in the famine-like cases, this provides some reason for thinking that such costs could also counteract the salient fact in the pond-like cases.

The moderate might try to respond by pointing to some feature (or set of features) that justifies treating the two sorts of cases differently. If this feature is peculiar to the pond-like cases, then we will want to say that it is this feature, in addition to the salient fact, that explains why one must act in Shallow Pond. Saying such things, however, is liable to sit uneasily with the earlier observation that other considerations often seem inappropriate as additions to an answer to the first component question. On the other hand, if we try justifying the differential treatment by pointing to additional features peculiar to the famine-like cases that counteract the salient fact in them, then this will likely sit uneasily with the fitting answer to the famine question—that it is the costs of repeated rescue that permit our inaction. If, moreover, the moderate's resistance to the demanding conclusion is indeed motivated by the threat of the high costs of repeated rescue then pointing to other features will look suspiciously like a rationalization.

So there are two ways in which our answers to the component questions increase the difficulty of the differentiation question. First, taken individually, the intuitively fitting answers to the component questions suggest, if anything, that the two cases should

be treated alike. Second, attempts to identify features that could provide an answer to the differentiation question threaten to run suspiciously afoul of the intuitively fitting answers to its component questions.

3.1 The promising general shape of a moderate solution

The moderate's most promising strategy, I take it, is to craft a proposal that explains why the two sorts of cases should nevertheless be treated differently, but to do this in a way that minimizes the conflict with the fitting answers to the two component questions. Following these guidelines, we would look for a moderate solution that has two features. First, it will give a central place to the threat of the high costs of repeated rescue in justifying our resistance to the demanding conclusion. Second, it will explain why we should regard the salient fact as providing decisive reason to act in the pond-like cases.

3.2 Recasting the challenge

Assuming this to be the most promising shape of a moderate solution, we can now recast the challenge facing the moderate into a problem with two interrelated parts. The first part of the challenge is to explain how it is that high costs of repeated rescue can counteract, at least in Distant Rescue, the salient fact. How can it be that the costs of having to help many imperiled people mitigate my duty to rescuing one person if that individual rescue can be performed at a very low cost to myself? To insist that my duties decrease as the costs of compliance increase seems to imply a strange relation between the duty of helping others and the needs of others: my duties decrease as the needs

increase. It is also, of course, to insist that I am morally permitted to allow a person to perish for want of the ten dollars that I could easily give. So, contrary to what many assume, there is no general moral duty of "easy rescue."[60] If moral requirements are to be understood as requirements to uphold minimal standards of decency in our relationships to others, this has surprising implications about how minimal those standards can be. If morally permissible behavior is that behavior which we can justify to everyone else, it is surprising to learn that we can justify to the perishing person our informed refusal to provide the ten dollars that would rescue him.

If the standards of minimal decency or for the justification of our behavior to others are so surprisingly minimal, we might wonder in what other ways we could be similarly surprised. This question is made acute by the fact that the moderate's strategy simply appeals to the threateningly high costs of repeated rescue, or the demandingness of a moral principle that requires one to respond to the salient fact in cases like Distant Rescue. If such costs of compliance can so affect moral responsibility, then we expect them to affect other moral issues as well (beyond the narrow range of cases involving rescue, beneficence, or even the so-called positive duties). If it becomes similarly onerous to avoid harming distant others, might we come to also acknowledge shockingly permissive moral rules about inflicting harm on innocent people?

The second part of the challenge is roughly this: if our moral duties of rescuing others are relaxed because of the high costs of repeated rescue, then what is to ensure that

[60] Henry Sidgwick claims that "all admit that we have a general duty of rendering services to our fellow-men and especially to those who are in special need, and that we are bound to make sacrifices for them, when the benefit that we thereby confer very decidedly outweighs the loss to ourselves" (1981: 348-49, cf. 261-62). Rawls offering this as an example of a "natural duty": "the duty of helping another when he is in need or jeopardy, provided that one can do so without excessive risk or loss to oneself…." (1999a: 97). The same is explicitly affirmed by many others (Arthur 1977: 48-49; 2002: 535; Smith 1990: 23-24). Michael Walzer (1983: 33), Thomas Hill (1997: 62), and Barbara Herman (2001: 227) note the general agreement on this point.

one has a duty to perform the rescue in the pond-like cases? If the high costs of repeated rescue can counteract the salient fact, what can possibly prevent them from so functioning in the pond-like cases (especially in the case in which Distant Rescue is conjoined with Shallow Pond)? However we answer this question, it will be important to bear in mind the observations about the seemingly fitting answers to the simple question of why one should rescue the imperiled person in Shallow Pond. When we point to some feature peculiar to the pond-like cases to explain why the consideration of the high costs doesn't counteract the salient fact, this explanation needs to somehow accommodate the fact that considerations beyond the salient fact are often intuitively inappropriate to various kinds of deliberations about such cases.

My strategy for dealing with the two-part challenge is first to identify in each of them a bullet that the moderate should bite. This done, the next steps will be relatively clear: seek to mitigate the damage. These steps, however, won't remain purely defensive. They also direct us further towards a promising positive account of rescue duties and of the relevance of costs of repeated rescue.

3.3 Biting the first bullet

Since the second part of the challenge requires the moderate to explain the duty of the rescue in the pond-like cases in terms of features that set them apart from the famine-like cases, the moderate's proposal is bound to sit uneasily with what may seem like an intuitively fitting answer to the question of why one must perform the rescue in Shallow Pond. Here, then, is a bullet that the moderate must bite. The moderate must concede that there is a sense in which it is perfectly correct to list factors other than the salient fact

when answering the question of why one must perform the rescue in the pond-like cases.[61] Depending on what these further factors turn out to be, it is at least possible that they won't strike us as so intuitively inappropriate. The moderate might try to select further features for differentiating the pond-like cases with an eye towards minimizing the intuitive inappropriateness of the answer. My own attempts to implement this strategy suggest to me that it doesn't provide adequate guidance.[62]

We might also try to lessen the uneasiness by providing an account of what makes some answers to the first component question seem more fitting than others. To do this,

---

[61] An implication here is that the refusal to rescue the child in Shallow Pond isn't (necessarily) wrong because it violates the rescue principle, as I tentatively suggested at the end of ch. 3, sect. 2.1 (page 48). That is, the threat of excessively high cumulative costs of rescue, posed by the threat of famine-like cases, may itself generally counteract the presumptive duty created by the salient fact.

[62] Attempts to implement this strategy lead me to think that it is not very promising. The idea is that some features beyond the salient fact will seem intuitively more inappropriate than others. Consider a variety of answers to the first component question ("Why should I rescue this drowning person?," where "this person" refers to the imperiled person in Shallow Pond):

1. "Because I can easily rescue her *and she is right in front of me*."
2. "Because I can easily rescue her *and I can hear her screams for help*."
3. "Because I can easily rescue her *and I've yet to do very much on behalf of others."*
4. "Because I can easily rescue her *and it would help sustain a valuable community practice."*

Answers 1 and 2 are intuitively more acceptable than answers 3 and 4 even if the further features identified in 1 and 2 do not, upon reflection, seem to be features that could carry much normative weight (mere physical proximity to the need; the salience of the need to the potential rescuer). By that way of scoring the features, answers 3 and 4 should fare better.

What makes answers 1 and 2 seem intuitively more appropriate, however, plausibly is not that they appeal to features that can themselves account for what is normatively special about the pond-like cases. Rather, the features appealed to in answers 1 and 2 may instead express sensibilities that are themselves made appropriate by other consideration that do the real normative work. Such answers may also be appropriate as appeals to shared sensibilities, or as appeals to common ground. Understood in this way, the answers function like a rhetorical question: "Isn't it also just obvious to you that, in a pond-like case, with such and such features, one must perform the rescue?" What the moderate needs to provide, however, is an account of what could ensure that the sensibilities of the kind triggered by the pond-like cases (physical proximity, the auditory salience of the need, etc.) are morally appropriate sensibilities to possess. The features that should trigger our responses of rescue in the pond-like cases may not of themselves reveal the appropriateness of their status as triggers. (These points do not, I think, undermine the relevance of the broader intuitively fitting responses to the first component question to the differentiation challenge. The moderate's solution to that challenge will deviate from commonsense to the extent that it ultimately fails to accommodate the intuitively fitting answers to the first component question.)

we could try to give an independent account of the special norms of practical deliberation that would explain why only certain sorts of considerations are appropriate in our deliberations about Shallow Pond. We might appeal to a two-level theory whereby the principles that should guide deliberation are different, though perhaps derivable from, the principles or values that determine the moral status of actions.[63]

Alternatively, we might try to justify the special status of pond-like cases in terms of a value whose realization necessarily depends upon our treating the salient fact as a decisive reason to act in such cases. Just as the value of my friendship with another person might be plausibly thought to depend on the way that I treat the interests of my friend—in my actions as well as in my deliberations—there might well be some valuable "relationship"[64] we have towards anyone in our immediate physical proximity (for example)[65] that likewise depends upon our granting to their interests a special kind of

[63] From a consequentialist perspective, and given our practical limitations, the optimal rules for deliberation may not always lead to optimal choices. As such, consequentialist standards for appropriate deliberation (or habit formation) may diverge from consequentialist standards for right action. At the level of the actions, however, such an analysis fails to distinguish failures to act in Shallow Pond from failures to act in Distant Rescue. Recognizing this creates a conundrum for the agent in any particular instance of deliberation about what to do in Distant Rescue. Perhaps it helps that the consequentialist can still know that his failure to act in Distant Rescue might plausibly be, from a consequentialist perspective, much less blameworthy, taking blameworthiness to concern the appropriateness of sanctions or punishments (Arneson 2004: 51-56). However, for the agent who is deliberating about what to do (as opposed to deliberating how to evaluate an action already accomplished, or about which habits to break or cultivate) the difficulty persists: for such deliberation the category of blameworthiness is out-of-place when the optimal choice about what to do is in this instance so clear.

[64] Since it is also the case that the imperiled person in Shallow Pond is a complete stranger to me, the notion of "relationship" will have to be somewhat broad. What we might look for is not necessarily a relationship that resembles prototypical "special relationships" (familial relationships, contractual relationships, and friendships, e.g.,) but rather a relationship whose value is similarly recognizable and important relative to the costs that respecting those values entails (e.g., costs to the range of our permissible options, and costs concerning the distribution of our concern and care for other people and for valuable projects).

[65] It would be its own project to identify the precise conditions of this relationship. For now, to simplify, I will speak of it as a relationship based on physical proximity to the person in need. If, however, intuitive moral judgments about cases are a reliable guide to the features of this relationship, there are interesting nuances to how proximity matters (while physical proximity seems to matter, what matters is not necessarily the proximity of the potential benefactor to the potential beneficiary) (Kamm 2007: 345-97).

unquestioned decisiveness in our deliberations about how to respond. In this way, the value of the proximity-sensitive relationship would imply both our reasons for performing the rescue and the appropriateness of taking the salient fact as providing decisive reason for doing so. Since this kind of ready explanation of the seemingly fitting answer to the pond question isn't available under all proposals for why one must perform the rescue in Shallow Pond, it suggests a way in which a special relationships proposal is distinctively promising.

If we account for the duty to rescue in the pond-like cases in terms of standards essential to a valuable relationship we share (at least) with whoever happens to be nearby, then this will also provide resources for dealing with the first part of the challenge. That part of the challenge concerns the counterintuitive permissiveness that the moderate's proposal threatens to imply. We must address the apparent fact that the requirements of morality sometimes are highly demanding.

To address this issue, we might try to set apart the kinds of cases in which extremely high moral demands are intuitively legitimate, showing that there is something characteristically special, and characteristically valuable, about following them. Alternatively, we might focus on the distinctive burdensomeness of the costs associated with the demanding conclusion in Distant Rescue. Both strategies hold promise and there is no reason why they shouldn't be combined.

Following the first strategy, we might try to separate out those highly demanding requirements associated with prototypical special relationships. The duties based in friendships and in family relationships can become highly demanding. Here, however, we might argue that when a moral duty is understood in terms of the standards essential

to a valuable relationship, it is intelligible why such duties are occasionally highly demanding. This is intelligible insofar as it is an important part of what makes such relationships valuable that the friend or the family member is prepared to meet certain needs even if the costs of doing so should become quite high in a perfect storm of unfortunate circumstances. It isn't only through blind tradition that marriage vows often include the phrase "in sickness and in health." If a child cannot expect her parents to remain at her side should she be afflicted by a debilitating disease, this diminishes the special value of that relationship, even if the child is perfectly healthy and no such unfortunate circumstances ever arise. Our most valuable friendships are not fair-weather friendships, even in fair weather.

Measured ex ante, or before the actual occurrence of the unfortunate events can be anticipated, the costs of participating in such relationships are unobjectionable in light of the value of those relationships.[66] If accepting those costs is essential to realizing the

[66] In appealing to the relevance of ex ante measurements the moderate will want to provide a principled way of steering clear of the view that ex post measurements of costs are to be generally replaced by ex ante measurements. If costs are always to be assessed ex ante, then a utilitarian principle that would require the sacrifice of one person's basic needs, when this is a way to produce small benefits for every one of a great many individuals, might seem to generate only tolerable costs for everyone. The expected costs to any individual may be quite small, or they might be offset by expected benefits to the individual that widespread compliance would imply. What is central, however, to the proposal I am considering is the intuitive value of the interpersonal relationship that depends on the potential benefactor's ex ante acceptance of the relevant costs. In the utilitarian case, no such intuitively valuable interpersonal relationship is realized by the ex ante acceptance of the cost of sacrificing oneself for the trivial gains of many others, should the opportunity to do so arise.

The proposal I am suggesting also contrasts with the contractarian idea that the duty of rescue in the pond-like cases is to be understood as an agreement of mutual aid, where such an agreement is mutually beneficial when the costs and benefits are assessed ex ante, and perhaps under some suitable restrictions on the contracting parties' knowledge of particular facts about their abilities, needs, or positions. The proposal I have in mind doesn't depend on the rationality of a quid pro quo agreement with others in one's community. Like the value of a friendship, the value of the proximity-defined relationship, as I see it, is not fundamentally the instrumental value to individuals of standing in that relationship. While there is value *to individuals* of standing in these relationships, such value does not reduce to the value to individuals of better securing their own narrowly defined self-interests—like their interests in the rescue of themselves or their loved ones, should the need arise. (Likewise, in contrast to a utilitarian perspective, the value of the proximity-defined relationship also doesn't reduce to, or arise from, the value of promoting the general good.)

value of these relationships, then it is more intelligible why special relationships can, in unfortunate circumstances, generate highly demanding duties. A similar story might be told about the special relationship we share with anyone who is (for example) nearby, even if they are strangers. We might explore the importance of being able to count on one another to respond as one might expect in the pond-like cases. In this way, the moderate proposal might help to account for certain cases in which morality is highly demanding, including certain cases of rescuing a stranger.

Following the second strategy, we might identify ways in which the compliance costs associated with the demanding conclusion in Distant Rescue are distinctively onerous, placing burdens on us that differ in kind from those we are still pressured to accept as legitimate. When facing Distant Rescue, it is not as if the demanding conclusion simply prevents me from purchasing a movie ticket. Rather, it essentially prevents me from ever engaging in discretionary spending on my own non-altruistic projects, or from ever sustaining a personal financial state in which, should I lose ten dollars, it would only be a small loss to me. If moreover, I am ever given the opportunity to make another ten dollars with just a bit more effort, or time, or inconvenience, I must take that opportunity, as this extra income will mean one more person whom I can rescue at a small cost to me. So, even if living without cash feels like a tolerable option,[67] Heroes for Hamiltons might add an option of accepting, as rescue-sized donations, small

[67] As it did to one graduate student, described in an article entitled "I live without cash – and I manage just fine: Armed with a caravan, solar laptop and toothpaste made from washed-up cuttlefish bones, Mark Boyle gave up using cash" (http://www.guardian.co.uk/environment/green-living-blog/2009/oct/28/live-without-money).

pieces of my time, allowing me at any time to rescue another person by clocking another hour on a menial online task.

What strikes me as particularly burdensome about the demanding conclusion in Distant Rescue is the way in which compliance so oppressively limits the range of permissible options left open to me. While I still want to concede the legitimacy of other highly demanding duties of rescue, such duties typically still leave me with a range of attractive options and possibilities going forward. Consider a high-cost case mentioned by Murphy:

> For years some astronomers have been planning an observation that must take place at sea at a particular date and time; many hundreds of thousands of dollars have been spent in preparation. Just as they are about to take their readings, they receive a Mayday signal. If they rescue the sailors in distress, the project will come to nothing. But, as we intuitively think, rescue they must. (2000, 127)

While the costs are great, the only ones that are emphasized are sunk costs. So, while the required sacrifice may frustrate significant past expenditures, the limitations on future prospects are highly circumscribed. To any one individual, the greatest prospective cost may fall on some intensely curious astronomer whose highly specialized career depends on the project's success. However, even in the unlikely event that the project's failure requires her to abandon her passionate professional pursuit of astronomy, the options remaining open to her are not limited in any shockingly extreme way.[68]

Elizabeth Ashford, claiming, "Any plausible moral theory must hold that there are some situations in which agents face extreme moral demands," illustrates her point with a

[68] Hershenov (2002), discussing Unger's case of Bob's Bugatti (page 134, below) and variants, argues in a similar vein, but emphasizes the psychological differences between prospective costs and costs that merely render one's past sacrifices futile. Even if, however, someone idiosyncratically found it disproportionately difficult to render past sacrifices futile, I suspect that this shouldn't greatly alter the wrongness of that person's refusal to rescue in cases like Bob's Bugatti or the astronomer's. The normative fact is largely independent of the psychological one: regardless of what a person is inclined to feel, sunk costs are generally not comparable to prospective costs.

scenario "in which the only way of stopping billions of people suffering an agonizing death [is] by hacking off your left leg with a fairly blunt machete" (2003: 274). While the cost here is both extreme and prospective, it is still very different in kind from the prospective compliance costs of the demanding conclusion in Distant Rescue. Once past the leg removal I am free to proceed, on one leg, however I choose.

The way that the demanding conclusion in Distant Rescue severely limits options cannot of itself, however, entirely justify our resistance to that conclusion. Plausibly, there are moral requirements, and even rescue requirements, that place similarly severe and ongoing limitations on our options. I might encounter a child drowning in a shallow pond in the midst of my escape from a Siberian labor camp, to which I have been wrongfully sentenced for life; I escape with my freedom if and only if I run past the drowning child. Alternatively, and possibly more convincing on the general point, I might escape with my freedom if and only if I torture and kill an innocent hostage.[69]

So while I think it is helpful to bear in mind the distinctive kind of cost associated with the demanding conclusion, this can only provide a partial explanation of the moderate's resistance to that demand. We will likely also need the support of the first strategy, showing, perhaps, that there is something characteristically valuable about accepting the extreme costs of the demands that the moderate is still pressured to accept as legitimate.

---

[69] The demanding conclusion in Distant Rescue is perhaps made more oppressive also by the *manner* in which it severely limits my options. Cullity mentions the "kind of self-surveillance that would have to characterize the life of anyone who was aiming to follow" the demanding conclusion in Distant Rescue (2003: 407). The duty to voluntarily hand over any discretionary income for the sake of morality makes demands of my will that aren't mirrored in any externally-imposed coercion I experience in the labor camp. I don't, however, believe that these additions can carry us all the way to a sufficiently principled distinction (Heroes for Hamiltons might offer the service of external coercion if I sign my name to a suitable contract, enforcement of which is provided pro bono by the Zetas drug cartel).

3.4 Biting the second bullet

However we characterize the burdensomeness of the high costs of repeated rescue, so long as we claim that these can counteract a moral requirement to rescue another person at a very small cost, we may have to concede that the threat of similarly high cumulative costs elsewhere can likewise counteract other elements of what is commonly assumed to be a fixed moral standard, including the so-called negative duties.[70] The alternative is to advocate a dubious exception for rescue obligations. Let us therefore bite a second bullet.

One way to minimize damage is to tell a plausible general story about how the moral standard for our interpersonal relationships can be sensitive to the anticipated cumulative costs to individuals of having to comply with such standards. In the previous chapter I argued that we might reasonably think that if I can easily rescue another person from serious peril at little cost to myself, then, by default, I am morally required to do so. This may seem like a standard of minimal decency in our relationships to other people as such, and we might think of moral wrongdoing generally as violations against such standards of minimal decency.

While we might assume that such standards of minimal decency are fixed regardless of great increases (or decreases) in expected costs of compliance, this assumption isn't necessary. We might alternatively suppose that the standards of minimal

[70] Robert Goodin writes that the "relative supremacy of negative over positive duties has long been established. Among the 'absolute duties…such as oblige all men antecedently to any human institution,' Pufendorf ranked 'first and noblest' the requirement 'that no man hurt another.' Besides the importance of the duty to maintaining 'the life of society,' Pufendorf writes that '[i]t is at the same time the most easy [duty] to be performed, consisting for the most part purely in a negative abstinence from acting'" (1985: 19, citing Pufendorf 1717, bk. 1, chap. 1, sec. 1, and bk. 3, chap. 1, sec. 1).

decency can be affected by significant changes in anticipated cumulative costs of compliance.[71] If we grant this supposition, then we can consider two ways of reducing the costs of compliance through adjustments to the standards of minimal decency. The first way is to weaken the content of such standards, so that they require of us even less in our relationships to everyone. The second way is to limit the scope of such standards, so that instead of requiring us to fulfill their requirements in our relationships to every person, they only require it with regard to some. When the standards of minimal decency are limited in this second way, they create a kind of special relationship between oneself and those within the newly limited scope of that content-preserving standard of minimal decency.[72]

---

[71] The flexibility of the moral standard of minimal decency might instead be characterized by a more complex inflexible moral standard, one that might be expressed with conditional clauses that render it sensitive to the same changes in anticipated cumulative costs of compliance. An unchanging conditioned standard could be extensionally equivalent to the adjustable standards that I describe in that, in the same circumstances, both always require the same actions. The unchanging conditioned standard might better resemble deontological principles with appended "catastrophe clauses" permitting, for example, someone to lie or to kill the innocent if doing so is the only way to prevent a nuclear war (Hill 1992: 199-200).

[72] These two different ways of reducing the costs of compliance track, I think, a commonly used distinction between special duties and general ones. There is, however, a way in which all such distinctions seem to collapse. Any special duty that relies on, for example, a special relationship, can be re-described as a general duty that (like any duty) only applies in specific circumstances. The specific circumstances in which the so-called special duty applies are simply those that constitute the so-called special relationship. Any special duty I owe to my child might therefore be re-described as a (general) duty I owe to anyone who happens to occupy the particular circumstances that my child happens to occupy in relation to me. Similarly, it should be possible to re-describe restrictions on a moral principle's scope as restrictions on a moral principle's content. My justification for using the distinction despite the possibility of collapse is that the idea of a special *relationship*—as opposed to simply special circumstances that trigger a duty—seems to convey an intuitively meaningful category of circumstances that are normatively privileged on account of the valuable interpersonal "relationship" we take such circumstances to constitute. The notion of "relationship" here is not as loose as the trivial one in which any one thing is related to any other in infinite ways, but it is also not as strict as, for example, Scheffler's. For Scheffler, the valuable special relationships that generate practical reasons "of partiality" are "ongoing bonds between individuals who have a shared history that usually includes patterns of engagement and forms of mutual familiarity, attachment, and regard developed over time" (2010: 115, 124). As I see it, what is relevant is not so much whether the relevant proximity-sensitive relationship contains the features common to prototypical special relationships that most clearly generate special duties (familial relationships and friendships, e.g.), but whether or not the value of the relationship is *similarly* recognizable and also important relative to the sorts of costs that respecting those values entails (costs, e.g., to the range of our permissible options and to our "moral autonomy," and costs concerning the distribution of our concern and care for other people). What

To illustrate this idea we might imagine a small and completely isolated island community whose culturally homogenous members are prosperous and secure. The standards of minimal decency in this world may be much higher than those of a less flourishing world. If any islander should happen to fall into the relatively dire straits of running short of good wine, minimal decency alone may require all other members to rescue the poor fellow by opening their own private cellars to her. In such a world, the expected costs of compliance of the standard would still be quite low, and it might even be gratuitously mean to reject such a standard, to which all may have come to appreciate and expect of one another. In such a world, this standard of minimal decency applies generally; it applies independently of any "special relationship" between particular persons in that world.

Suppose, however, that this island community suddenly makes contact with a billion less prosperous wine lovers heretofore not known to exist. The expected costs of compliance with the former standard of minimal decency suddenly skyrockets. This plausibly licenses an adjustment to that standard. Under certain circumstances, it might be appropriate to maintain the original standard, but simply to limit its scope (perhaps to include only one's fellow islanders, or maybe just one's immediate neighbors). In such a

---

we need, for my purposes, is something that is recognizably a relationship between persons, the value of which can make it more intelligible why the circumstances in question give rise to practical reasons.

A somewhat related distinction that avoids such collapse is the distinction between agent-relative and agent-neutral reasons/duties. Special-relationship-generated duties might be understood as agent-relative duties, where an agent-relative duty would be any duty that makes some kind of essential reference to the particular agent to whom it applies (Nagel 1978: ch. 10): it is only *because agent **A** stands in such and such circumstances* that ***A*** has such and such duty. Drawing the distinction between special-relationship-generated duties and general duties this way might in other ways be useful, but it seems to come with at least two costs: the interpersonal component of the duty becomes inessential, and the discussion is rendered into terms the significance of which isn't as intuitively accessible as "special relationship." We could avoid the first cost by characterizing special-relationship-generated duties as agent-relative duties generated by the interpersonal relationship I stand in with the person to whom the duty is owed. Doing so, however, seems to sideline the contribution of agent-relativity, placing the idea of the relationship again at the forefront.

way, what had formerly been a general standard of minimal decency might now become a special one, helping to define a kind of special relationship.

If we accept that significant increases in what would otherwise be the expected costs of compliance can justify cost-reducing adjustments to the standards of minimal decency, then we might further speculate about which of these two kinds of adjustments is more appropriate. Recall that the standard at issue is that of rescuing another person when this can be done at a small cost to oneself. If we have conceded that the extreme costs of compliance license alterations to this standard of minimal decency, then we have already eliminated the claim that we are required to fulfill *that* requirement in our relationships to every other person. Now the question becomes whether the costs of compliance should be reduced through a weakening of that standard with regard to all persons, or if instead the scope of that requirement should simply be limited. How we decide this question might plausibly depend on a variety of factors, including the value of maintaining that original standard with regard to at least some people, the available options for how we might limit the scope in a way that might secure that value, and the costs of treating people unequally (including the disrespect that the differential treatment might express towards those standing outside the potential special relationship).[73] We might try to formulate higher-order principles that govern the formation of special relationships in the face of changing circumstances.

[73] I stress the last consideration in this list so as to preempt the charge, which is sometimes made against those advocating for a moderate position (or the current "moral status quo"), that the kind of justification I am offering would have provided the antebellum slave owner or a racist with a straightforward defense of his/her differential treatment of those outside of his/her group (James 1982: 4; Kagan 1989: 14; Unger 1996: 19-20).

Could alternative changes in circumstances make it the case that seemingly fixed deontological prohibitions lose their status in the general standard of minimal decency? An affirmative answer to this question is that second bullet that the moderate will have to bite. Here, however, I am not suggesting that all moral principles should be equally sensitive to the cumulative costs of compliance—as if the requirement against killing others should, in the face of threateningly high cumulative costs of compliance, be compromised as readily as a requirement to help others further their non-vital interests. Some standards are more important in our interpersonal relationships than others.[74] I am instead saying that when we resist the demanding conclusion by claiming that the moral standards of minimal decency are not fixed, but are rather open to adjustment as expected cumulative costs of compliance change, then we cannot simply assume that this sensitivity applies only to the part of the moral standard that deals with the rescue of strangers.

The implication that there is no general duty of easy rescue—that is, no general duty to rescue someone from serious peril when doing so can be done at a very small cost to oneself—may be surprising largely because we didn't expect that circumstances would make the cumulative costs of complying with such a duty so shockingly high. This would fit the earlier discussion about a student's typical progression of thought when initially confronted by the famine-relief argument. If, then, it is initially so

[74] On this point, Mill's observation is apt:

> The moral rules that forbid mankind to hurt one another…are more vital to human well-being than any maxims, however important, that merely point out the best way of managing some aspect of human affairs….Their being observed is the only thing that preserves peace among human beings: if it weren't for the fact that obedience to them is the rule and disobedience the exception, everyone would see everyone else as an enemy against whom he must be perpetually guarding himself. (1993: ch. 5)

counterintuitive that there may be no general duty of easy rescue, this is because the circumstances created by Heroes for Hamiltons are so unexpected.

The moderate can tell a similar story about other standards of morality commonly assumed to be fixed, or which are commonly assumed to be a part of moral standards of minimal decency in our interpersonal relations. We assume such standards are fixed because we don't expect the anticipated costs of compliance with them to similarly become shockingly great. It is difficult to imagine, for example, a scenario in which a general prohibition against breaking a promise to avoid a trivial inconvenience could entail, through accumulation, extremely high anticipated costs of compliance to an individual. Promises after all are the sorts of things one can simply avoid making if the anticipated costs of having to keep them would be too high. If, however, unexpected circumstance do arise in which seemingly fixed deontological prohibitions come to have exceptionally high anticipated costs of compliance, perhaps this would make us rethink whether or not the prohibition should be part of the standard of minimal decency. What we should conclude would presumably depend not simply on the anticipated costs of compliance, but also upon the positive values associated with the prohibition in question.

While it is easy to imagine cases in which the costs to an individual of having to keep a promise or avoid theft or murder is exceedingly great on a particular occasion, it is difficult to imagine cases in which the costs of compliance with such standards resemble the sorts of costs associated with a general duty of easy rescue in the famine-like cases. It seems characteristic of negative duties, as opposed to positive ones like a duty of beneficence, that compliance with them still leaves people with plenty of freedom to pursue other interests. To find a comparable case in which the prohibition against killing

an innocent person is similarly constraining requires a fictitious scenario quite a bit more fanciful than Distant Rescue. Here is such a case:

> *Intergalactic Killer*: We suddenly realize that our universe is much stranger than we previously thought when NASA discovers the following series of startling truths:
>
> 1. Distant galaxies are full of humanoid creatures that perfectly resemble ourselves. Were we to meet them here on Earth, we would simply take them to be people from an unfamiliar country. Nothing about these humanoids otherwise distinguishes them from normal human beings.
>
> 2. Peculiar one-way causal laws connect every individual on Earth to huge populations of these distant humanoids. The content of these laws is this: over the course of an average lifetime, each of us on Earth causes the early deaths of hundreds of thousands of distant humanoid creatures. The way I cause such deaths is similar to the way everyone else on Earth does: for every day that I continue to live, the simple fact that I have lived another day causes the early death of 48 distant humanoids, at a rate of exactly two deaths per hour.
>
> 3. Every day I can avoid killing humanoids in only one of two ways: I can commit suicide, or I can spend the entire 24 hours climbing stairs. Every half-hour that I climb stairs is a half-hour that I continue to live without killing a humanoid.[75]

While I expect moral judgments will differ about what one must do in Intergalactic Killer, I wouldn't be surprised if many judged as supererogatory the endeavor to fully maximize the portion of one's life spent stair-climbing, and as permissible the decision to

[75] If negative duties are characterized as duties that merely require one to refrain from doing things, in contrast to positive duties requiring one to perform certain acts, then the avoidance of killing would, in this circumstance, count as a positive duty. To fit this characterization of negative duties, we can instead imagine that, in order to avoid killing the distant humanoids, one must refrain, for the entire 24-hour period, from moving one's limbs and opening one's eyes.

We might make Intergalactic Killer resemble Distant Rescue in other respects if we add further features. We could, for example, also stipulate that scientists discover an easy method for allowing one person to vicariously stair-climb for someone else. While a vicarious stair-climber continues to kill humanoids at the same rate as always, his/her stair-climbing prevents another person's life from killing. An innovative organization employs confirmed sociopaths (people who would not otherwise spend their days climbing stairs) to vicariously climb stairs on behalf of anyone who is willing to pay the considerable fee. Through this organization, you can now avoid killing humanoids if you are willing to spend your life either climbing stairs or earning enough money to continuously employ a vicarious stair-climber. Either way, at least for most of us, all leisure time or discretionary income would have to be sacrificed if one is to live and to avoid killing innocent others.

continue living largely as before alongside a serious commitment to a regular stair-climbing routine, and to give up jogging and the gym. Such a judgment would concede the absence of a general moral prohibition against killing an innocent person, even when such any individual killing could be avoided at only a modest cost to oneself (a mere half-hour of stair-climbing). I would, however, still expect most to maintain the conviction that it is impermissible to deliberately run over a stranger with one's truck, even if the cost of doing otherwise is unusually great (perhaps it will mean missing your very important flight).

3.5 Conclusions

I have argued in this chapter that the difficulty of the differentiation question can be appreciated through an examination of its two component questions. Viewing the difficulty in this way led me to suggest that the moderate's most promising path lay in accepting that natural idea that it is at least in part because of the costs of repeated rescue that I needn't treat the opportunity in Distant Rescue like the opportunity to rescue in Shallow Pond. To better accommodate the idea that the costs of repeated rescue can counteract the salient fact, I suggested how moral duties more generally are plausibly limited by compliance costs of the same sort, and that sometimes such limitations can create "special" duties which we only have in more circumscribed circumstances, or only towards select people. Such duties might be said to correspond to "special relationships."

The answer to why I must, nonetheless, continue to perform rescues in Shallow Pond remains underdeveloped. I concluded that the moderate must accept that the question of why I must rescue the person in Shallow Pond isn't adequately answered by

an appeal to the salient fact. Since, however, appending further considerations to the salient fact tends to make for awkward answers, the moderate's best strategy is to account for that uneasiness. This led to the thought that whatever it is that might explain my continuing duty to perform the rescue in Shallow Pond (in the face of Distant Rescue) might also explain why, in such cases, I should also take the salient fact as providing decisive reason to act. Since properly valuing special relationships often also involves special standards of deliberation, the idea of explaining the circumscribed duty to rescue in the pond-like cases in terms of something like a special relationship seemed attractive again.

What prominently remains, of course, is the task of explaining how pond-like cases can be intelligibly viewed as involving anything like a special relationship, lending, as it seems, such heightened significance to salient fact that there is in them a special duty to respond. In the next chapter I will approach this task; before doing so, however, I will try to clarify better the features of the duty to rescue in the pond-like cases, showing how these features comprise, in the context of the differentiation question, a neglected challenge.

CHAPTER FIVE:
THE STRINGENT DUTY TO RESCUE

1. A NEGLECTED QUESTION

Moderate responses to demandingness challenges often attempt to explain limits on what can be morally demanded of us. I argued in the last chapter that the moderate's most promising strategy in response to question why, in Distant Rescue, one need not keep responding to the salient fact is to propose that moral requirements generally are sensitive to the cumulative costs to individuals of compliance. To leave off here, however, would be at best to answer only half of the differentiation question.

A weakness of the most developed moderate responses to the famine-relief argument is that they do tend to leave off more or less here. Efforts are aimed at responding to the famine question while the pond question is comparatively neglected. Garrett Cullity's arguments primarily attempt to explain why the repeated costs of rescue limit our duties in the famine-like cases.[76] The task of explaining why we continue to

[76] Cullity has published a number of essays on this problem at least since "International Aid and the Scope of Kindness (1994), culminating in *The Moral Demands of Affluence* (2004). He reaffirms the structure of his central argument for resisting the demanding conclusion in "Demandingness and Arguments from Presuppositions" (2009).

have a duty in the pond-like cases, however, seems to arise as an afterthought. Cullity seeks to partially explain our intuitive moral judgments in the pond-like cases as judgments about character or blameworthiness, as opposed to judgments about the moral impermissibility of inaction. Since these explanations fail to account for the moral obligation of acting in Shallow Pond, Cullity then suggests that there might yet be further features of the pond-like cases that somehow create stricter obligations.[77] As a contrast to a case like Distant Rescue, however, Cullity doesn't offer any indication of what these further features might be (2004: 199-201).[78]

Similarly, while Liam Murphy's *Moral Demands in Nonideal Theory* offers a detailed justification for limiting the general demands of beneficence, he offers only a brief discussion about the pond-like cases. Like Cullity, he suggests that our intuitions

---

[77] At one point in his argument Cullity provides a principled reason for distinguishing pond-like cases from *certain sorts* of famine-like cases in which the charitable organizations operate in such a way that no individual contribution has any chance of making a significant difference in any needy person's life. For Cullity, one's obligation to the distant needy in such cases is an obligation to do one's fair share of a collective action (2004: 60-65, 2000: 14 ff). Pond-like cases are therefore importantly different because they provide one with an opportunity to perform an action that will of itself rescue a person, and this explains, for example, why, in pond-like cases, performing the rescue takes precedence over preventing threats (2004: 66-67). Believing that that international aid organizations function in such a way that individual contributions don't make any difference (or any significant difference) for anyone (2004: 58-60, 2000: 2-3, 1996: 53-55) (a point emphasized in Whelan 1991: 158-62), Cullity doesn't consider more idealized cases like Distant Rescue, or peculiar organizations like Heroes for Hamiltons.

One might try to argue for additional significance of this de facto difference between pond-like cases and famine-like cases by something like Violetta Igneski's suggestion that it would be morally wrong for aid agencies to set themselves up like Heroes for Hamiltons (2006: 452 ff). Alternatively, following Onora O'Neill, one might try to argue that if cases like Distant Rescue "are *very* unlikely, we don't need a moral theory" to deal with the problems they pose. While there may be other ways of supporting this conclusion, O'Neill's own reasoning is narrowly pragmatic: "There is little point in strenuous thought about unlikely problems when there are so many problems that we know will arise and require decisions…." (1986b: 300). Here, as well as in *Faces of Hunger* (1986a), O'Neill's efforts are primarily aimed at discovering, for the sake of implementing, general policy guidelines for individuals and institutions in addressing global needs as they typically exist.

[78] Cullity accommodates my observations about the second component question without fully addressing the implications this might have for the first component question (and, of course, without fully addressing the composite question). Cohen (2005) exemplifies a response to the famine-relief argument that accommodates my observations about the first component question without accommodating my observations about the second (and, so, without providing a moderate answer to the differentiation question).

about such cases might be understood in terms of blameworthiness judgments directed towards the agent's character, as opposed to judgments about the moral impermissibility of inaction. Unlike Cullity, Murphy expresses skepticism that any special duty applies in the pond-like cases on account of peculiar features of such scenarios. Murphy suggests instead that surprising revisions may have to be made to how we might otherwise evaluate certain pond-like cases (2000: 127-33).[79]

Such treatments don't do justice to the challenge posed by the pond question. To introduce this particular challenge, recall some previous observations. I argued that two considerations must be at the forefront of any moderate solution to the differentiation question. One is the importance of responding to the vital needs of others; the other is the cumulative cost, to an individual, of having to meet the needs of others. It is tempting to conclude that a fitting way to honor these two considerations in Distant Rescue would be to adopt the following personal policy: allocate a generous percentage of one's personal resources to the work of rescuing others and then maximize the number of people whom one rescues with those resources. After all, if my duty to imperiled strangers is limited by the costs of compliance, then we might reasonably suppose that the needs of imperiled strangers can only legitimately claim some portion of my resources. In earmarking such resources ahead of time, I make a good-faith effort to set aside an amount that is safely beyond what the needs of imperiled strangers can legitimately claim. Having done this, it only seems reasonable to then direct these resources in the way that most efficiently addresses the needs. However, despite the apparent reasonableness of such a plan, it fails

[79] As we have seen, for Murphy, if it is less costly to rescue a person overseas, or if you have already done your fair share in our collective responsibility of beneficence, or if you are already very poor, then, counterintuitively, you may be under no obligation to rescue the nearby child from drowning in the shallow pond.

to respect the stringency, or the non-optional priority, of our duty to perform the rescue in the pond-like cases. A pond-like case could arise which requires me to deviate from my plan, forcing me to spend resources on the pond-like case beyond those I have allocated for meeting the needs of imperiled strangers. Since, moreover, pond-like cases can require me to perform higher cost rescues, they can also force me to deviate from my plan to maximize the efficiency of what I sacrifice on the behalf of others' needs.

To answer the differentiation question, then, the moderate must account for this relative stringency of the duty to rescue in the pond-like cases. To justify the stringency, the moderate seemingly must appeal to some third factor that sets the two types of cases apart. Since the factor must mark a difference between pond-like cases and famine-like cases we can call it a *differentiating factor*. The mystery is how the distinctive features of the duty to rescue in the pond-like cases can be explained in terms of any such factor, or set of such factors.

2. The Challenge of Stringency

2.1 Priority

When we note the stringency of the duty to rescue in the pond-like cases we are recognizing in them a feature more typically associated with deontological constraints. What is being constrained, or restricted, are certain kinds of seemingly reasonable, or attractive, trade-offs. I must not kill a passerby even if this means significant loss to myself, to others, to the preservation of biodiversity, and to the advance of the arts and sciences. Deontological constraints characteristically prohibit *value-increasing tradeoffs*,

whether the value is some supposedly overarching one, such as aggregate well-being, or a specific value that is seemingly driving the prohibition itself. Deontological constraints also characteristically prohibit *violation-decreasing tradeoffs*. These are the tradeoffs that would reduce the overall number of violations of the same constraints. I must not kill a passerby even if doing so would prevent five other people from killing five other passersby.[80] While it is the latter kind of tradeoff that gets associated with the so-called paradox of deontology, the former isn't less puzzling, or less in need of explanation. If there is something puzzling about the paradox of deontology, then it is presumably because the point of the constraint is taken to be the minimization of violations. Put another way, if the point of the constraint against killing is something other than the minimization of killings, then a prohibition against violating the constraint to decrease killings is no paradox. If there is something paradoxical, it is a constraint against tradeoffs that appear to be advantageous in terms of the very considerations that justify the constraint.

In what ways, then, do duties of rescue exhibit these paradox-inducing features? Performing a costly rescue of one person in a pond-like case may so deplete my set-aside charity fund that the remainder will finance the rescue of five fewer people. In such a case, however, I could typically just add more money to the fund; I can still rescue just as many distant strangers as I otherwise would have.

In a case like Peter Unger's Bob's Bugatti, however, the resemblance to the paradox of deontology grows. In Unger's case, 70-year-old Bob has invested nearly all of his assets into an uninsurable Bugatti automobile whose market value is expected to

[80] This is not to say that the constraint cannot be overridden if we greatly increase the numbers who would be killed. Perhaps it is permissible to kill a passerby if doing so is the only way to prevent five other people from killing five *billion* passersby.

increase dramatically over the next two years. Bob plans to then sell the car for several million dollars and enjoy a comfortable retirement thereafter. Here is the case.

> *Bob's Bugatti*. On a rural road near the garage where it's securely kept, Bob's gone for a careful drive in his Bugatti. At a certain point, he spies a shiny object. To inspect it, Bob parks his car in the only place from where, directly, he can proceed on foot for a close encounter, a parking place that's just ten yards beyond the end of a certain trolley track. As it develops, when Bob walks over to the shiny object, he finds it's a switch that can be set in two ways. And, as Bob observes, there's a trolley up the line that's barreling toward the switch's fork. As the shiny switch is set, the trolley will go down the fork's opposite side, not the branch leading to a spot near Bob's Bugatti. But, as Bob sees, on that side there's a young child trapped on the track. As he knows, Bob has two options: If he does nothing about the situation, the child will be killed…. If he changes the switch's setting, his second option, then, while nobody's killed, after rolling down the vacant branch and beyond that track's end, the trolley will totally destroy Bob's uninsurable Bugatti, wiping out his entire retirement fund. (Unger 1996: 136)

It is plausible to me that Bob is morally required to perform the rescue. Suppose, however, in what I will call the *augmented case of Bob's Bugatti*, that Bob's predicament occurs amidst the standing offer of Distant Rescue. Bob, then, might do a quick calculation. If he rescues the child, he will be able to rescue far fewer children through Heroes for Hamiltons. So Bob reasons with his own conscience: "Let me preserve my Bugatti and I will, in exchange, commit to donating $100,000 to Heroes for Hamiltons from its sale. In this way I will rescue many more people, since this is far more money I would otherwise be capable of donating." If Bob is nevertheless required to activate the switch and rescue the child, then a value-increasing trade-off is curiously prohibited. Rescuing the one stranger in the pond-like case requires forgoing his own opportunity to rescue many other strangers.

We might also imagine cases in which an agent can, by refraining from performing a pond-like case rescue, ensure that several other agents would perform pond-like case rescues that they otherwise wouldn't have. We can imagine encountering

something like Shallow Pond with this additional feature: several very similar pond-like scenarios are about to occur overseas, and only if I refrain from rescuing the drowning child in front of me will several other agents, far away, rescue the drowning children in front of them.[81] If I am still required to perform "my" rescue in this pond-like case, then a violation-decreasing tradeoff is also prohibited: I must perform the pond-like case rescue even if by doing otherwise I could decrease the number of cases in which people failed to perform such duties.[82]

The moderate isn't committed, as far as I can see, to any particular judgment in these particular high-cost rescue cases. The moderate might join those who deny the legitimacy of the constraints that generate paradoxes of deontology.[83] There will remain, however, some degree of that characteristic non-optional priority over seemingly reasonable alternatives that makes those constraints seem plausible for duties of rescue in

[81] Perhaps it is puzzling how my failure to perform a pond-like case rescue could result in the violation of five other agents of their pond-like case rescue duties far away (my causal role here may appear to stand in conflict with the requisite agency of these five others). Suppose therefore, that on any given day, hundreds of people encounter pond-like rescue scenarios, but that today, because of an unusually widespread malaise triggered by the defeat of the beloved underdogs in the World Cup, the number of people willing to violate their duties to perform such rescues is significantly higher. Fortunately, I have the opportunity to significantly counteract this malaise by ordering the global broadcast of a particularly inspiring Pepsi advertisement that is otherwise scheduled to air next week. By doing this today there is a statistical certainty that at least five people who would have refrained from performing a pond-like case rescue will instead perform their duties. Since the unusual malaise is temporary, airing the advertisement on any other day will have no similar effects. Unfortunately, in order to place my time-sensitive order to advance the broadcast, I must run past the drowning child in the nearby shallow pond.

[82] While I find it plausible that one is still required to perform the rescues in these cases, there are other similar cases in which it seems permissible to refrain. Suppose there are five people drowning in the shallow pond, and five onlookers, where each onlooker, including myself, is physically able to rescue only one child (imagine that the onlookers are evenly spaced around the same small pond, and there is a child drowning in the shallows next to each onlooker). Suppose further that the other onlookers will rescue children only if I refrain from rescuing. In this case I find it more plausible that it is permissible for me to refrain from rescuing. As with the other cases, however, this strikes me as the sort of intuition that I might be talked out of relatively easily.

[83] That said, I think that some cases like these, though perhaps not quite as extreme (e.g., perhaps the agent is not faced with a choice between saving just one nearby person or many thousands of distant ones) will have to be counted as pond-like cases if the moderate position is to avoid being deeply at odds with intuitive moral judgments about rescue cases.

the pond-like cases. I will call this feature *deontic priority*. The moderate should explain the deontic priority of rescue requirements in pond-like cases, even if she also tries to debunk it to some extent.

Providing the explanation is bound to be challenging. The paradox of deontology is said to be "formally puzzling," having an "apparent air of irrationality" (Nozick 1974: 30, Nagel 1986: 178, Scheffler 1994: ch. 4). There is no consensus on how these "paradoxes" are best understood or defended. Moreover, the proposals that are offered tend to be customized for principles that are crucially different from the duty to rescue in the pond-like cases. So, for example, any account of deontological constraints that places weight on the doing vs. allowing, or the harming vs. not aiding, distinction is a non-starter for our purposes. It is likewise difficult to see how accounts based on the inviolability of persons, upon the doctrine of double effect, or the Kantian restriction against treating others as mere means could be extended to duties of rescue.[84] In many cases, the associations of the paradigmatic deontological constraints with negative duties seem to have prompted accounts that are unlikely to extend to positive requirements of rescue.

I have argued, then, that the difficulty facing the moderate resembles a puzzle, the so-called paradox of deontology. Given our lack of any suitable explanation of these features of the pond-like cases, the prospects of finding a satisfying moderate response to the differentiation question may seem to be weakened. More optimistically, however, we might also say that we have come upon an opportunity to approach a well-known puzzle

[84] I don't mean to suggest here that Kantians lack other resources for explaining deontological constraints. Here I mention only one aspect of Kant's second formulation of his Categorical Imperative. More promising might be a different aspect of that same formulation, requiring us to treat people, or the humanity in the person of any other, always as an end.

of morality from a new and distinctive angle, one that might shed more light on the problem, or at least force us to look beyond the various commonly suggested proposals.

2.2 Priority and special duties

Towards the end of the last chapter, I suggested that it might be helpful to try to understand the duty to rescue in the pond-like cases as a kind of special duty. This idea initially arose when seeking to harmonize suitable answers to the differentiation question's component parts.[85] The current challenge of deontic priority gives us another reason to explore this idea. Special duties, like duties of rescue in the pond-like case, sometimes also exhibit deontic priority. Generally, I must keep my promise even if breaking it would allow me to prevent five others from breaking their comparably significant promises. I ought not to betray my friend even if through so doing I might somehow prevent five strangers from betraying their friends.

Unfortunately, it is also unclear how special duties come to have "paradox" generating deontic priority. It is doubtful whether we even have a compelling general understanding of special duties and of why they command our respect. This shows the lack of substance in the current suggestion that the pond-like cases somehow trigger special duties. At best we have proposed an association with the sorts of duties that moral philosophers classify as special duties. This, however, reveals a further concern. Characteristically, the duties considered special can be traced to some significant past transaction or some special relationship between the duty-bearer and the one to whom the duty is owed. In the paradigmatic pond-like case, however, I have a duty to rescue a complete stranger, and I cannot trace my duty to any special role or responsibility that I

[85] This discussion is in the latter part of ch. 4, sect. 3.3 (pages 114 ff).

have voluntarily assumed. So the duty of rescue in the pond-like cases is not an easy fit. Simply to insist that it is a *special* special duty is unsatisfying.

To build a case that pond-like scenarios trigger special duties of rescue, we need a conducive account of special duties. Voluntarist accounts aren't promising, since it is doubtful that I have voluntarily assumed any special responsibility for the imperiled person in Shallow Pond. It is helpful, therefore, that at least some special duties are plausibly independent of the agent's voluntary assumption of responsibilities.[86] Of the non-voluntarist accounts, we might discount those that assume a kind of value monism, attempting to account for all moral duties generally, or even of all special duties, by appealing to some single kind of value. Such accounts hold little promise for maintaining the desired contrast between highly idealized pond-like cases and famine-like cases. What could plausibly be the single value that accounts not only for special duties, but also for the desired contrast between of all of the strategically idealized pond-like cases and famine-like cases?

Offering a non-voluntarist, non-monistic account of special duties, Samuel Scheffler considers "relationship-dependent" reasons, suggesting that these arise not because one has voluntarily accepted the terms of the relationship, but rather because of the value of the relationship itself (1997: 196-98, 2010: 110). Scheffler argues that, as a further consequence of the relationship in question, these special reasons become special duties with deontic features.

[86] This is not to deny striking exceptions in, for example, Hart (1955: 185), Rawls (1999: sections 18, 52), and Thomson (1971: 65). Robert Goodin argues that voluntarism (the idea that special duties are voluntarily assumed responsibilities) fails to accommodate our "preanalytic" judgments about the special duties we have (1985: chapters 3-4). Scheffler addresses the apparent motivations that drive voluntarism (1997, 2003).

Rejecting any monistic restriction on what the value is, Scheffler's proposal allows us to think broadly about what could be generating special reasons for rescue in the pond-like cases. While I may have no relationship, in Scheffler's restricted sense, with the imperiled person in Shallow Pond, there may well be something distinctively valuable in that people relate, and are expected to relate, in the sorts of emergency scenarios that pond-like cases comprise. The very norms people ordinarily accept with regard to the pond-like scenarios could itself comprise a social relationship, the value of which we might appreciate by considering its absence. This is Richard Miller's observation:

> [T]he expectation that others who encounter us would help us if we needed to be rescued from imminent peril makes us much less alone, much more at home in our social world. Even if I were guaranteed not to need help in emergencies from mere passersby—say, because official emergency services were so wonderfully effective—I would be profoundly deprived of fellowship if those whom I encountered typically had no such inclination to help me if need be. (We find it chilling if someone 'looks straight through us,' even if we know this person is intensely active in relieving neediness worldwide.) (Miller 2004: 376)

So, between people who encounter one another, or who simply might encounter one another, there is something like a valuable relationship that we can call the *assumed practice of mutual aid*. This practice is the public's general respect for the norms surrounding the pond-like cases, which require us, in particular, to perform the rescue should such scenarios arise. The value of this relationship, which I will call the *Millerian value*, might give rise to special, relationship-dependent reasons. If Scheffler can then show us how reasons of this sort can give rise to special duties, we might also have a way of explaining how special duties arise in the pond-like cases. We would then have identified a recognizable value which is distinct from the importance of rescuing others generally, which distinguishes pond-like cases from famine-like cases, and which

generates duties that belong to a category for which deontic priority is not uncommon. The differentiating factor would be those special reasons generated by the value of the assumed practice of mutual aid.

That, at least, is the hopeful idea. The task of developing it, however, faces several problems. Part of what makes me hopeful about this approach is Scheffler's confidence that relationship-dependent reasons can create special duties, where these in turn might exhibit deontic priority. The first problem, however, is that Scheffler himself would reject my appropriation of his proposal. As he sees it, the only kinds of relationships that give rise to relationship-dependent reasons are "ongoing bonds between individuals who have a shared history that usually includes patterns of engagement and forms of mutual familiarity, attachment, and regard developed over time" (2010: 115). Whatever relationship binds me to any stranger I might encounter in Shallow Pond, it isn't one that fit Scheffler's criterion. Scheffler argues that it would undermine any explanatory role for the relationship if we understood relationships otherwise (123).

A second concern is that it is doubtful that Scheffler's own account succeeds in explaining how relationship-dependent reasons can ever amount to requirements of the kind we need to explain—requirements that posses their distinctive kind of stringency or priority. The puzzle of deontic priority is tricky, and I will later argue that Scheffler's proposals are inadequate.

2.3 Significance

Lacking the help we need from Scheffler, we are left with a considerable gap between the Millerian value and the deontic priority of the duty of rescue in the pond-like

cases. The problem, however, is still greater. Deontic priority is only one of the features that a simple appeal to the Millerian value leaves unexplained. Attributing priority of any kind to a duty is bound to seem inappropriate, and perhaps strangely legalistic, if violations of that duty do not strike us as correspondingly significant, failing to resonate with attitudes of remorse, resentment, or personal grievance.

That rescue duties in the pond-like cases do have a great deal of significance isn't doubted. My refusal to rescue the drowning child in Shallow Pond is horrific. What is doubted is whether this kind of significance can be justified if we account for the duty in terms of anything like the Millerian value. The Millerian value is the value of a social practice. It is perhaps a kind of public good. Speaking for myself, it is difficult to feel deep moral remorse at my failure to support an important public good, particularly when I know that my isolated failure does little harm to the public good itself.[87] This remains true even if I ascribe great importance to that public good. Towards failures of this kind, I can often clear my own conscience just by resolving to do better next time, or by doing something to compensate for any damage I have already done to the public good (trivial as that damage, and the compensating effects of my efforts, may be). So such an explanation of the duty seems to be phenomenologically amiss.

2.4 Direction

In accounting for the duty to rescue in Shallow Pond in terms of my responsibility for sustaining a public good, what seems amiss is that, under such an account, the value at stake is impersonal. Missing is any appreciation of the personal stake that another person has in my fulfillment of the duty, or of how my violation of the duty would let

[87] Here I draw ideas from Scanlon's remarks on what he has called "the remorse test" (2007: 8-9).

down the particular person whom I refrain from rescuing. So there appears to be a connection between the significance of the duty and its directedness. My sense of deep remorse at a refusal to rescue in Shallow Pond seems necessarily to involve the appreciation of how I have failed the particular person who has perished for want of my rescue (along with that person's family, friends, and dependents). It seems to me that the only way to have the full and seemingly appropriate remorse is to include in the wrongdoing a directional aspect.

As with priority and significance, this directional aspect is also characteristic of deontological constraints. Such constraints often involve corresponding rights or claims held by others. The constraint that prohibits me from deliberately sneezing on a passerby corresponds to the passerby's claim, upon me, not to sneeze upon her. Refraining from this is something that I owe to the passerby. So, if I do deliberately sneeze on her, then I haven't simply acted wrongly, I have wronged her. This wrongdoing has a particular victim, and she would have a grievance that she could make on her own behalf against me. All of this characterizes what I will call the *deontic direction* of deontological constraints.

It should be clear that duties of rescue in the pond-like cases have deontic direction. In Shallow Pond, the imperiled person has a claim to my assistance; I have a duty or obligation to the imperiled person; or, if I refrain from performing the rescue, I haven't just behaved poorly, done wrong, or evinced moral defect in my character, I have, in addition, wronged the imperiled person, and this person would then have a legitimate grievance against me (Feinberg 1987: 150).[88]

[88] Some resist certain of these propositions. Sometimes this is because of a sweeping rejecting of any duty to perform the rescue in Shallow Pond (Lomasky 1995). Others may simply reject the appropriateness of

While the deontic direction of such rescue duties is well attested, it is again unclear how we can account for this feature. If the Millerian value is the value of a public good, then the reasons it generates might simply be those that arise from my responsibility to support the public good. A duty of this kind, however, seems to be owed to the larger community, and not perhaps to the imperiled person.[89] When I refuse to rescue the imperiled person, such an account may strangely explain my wrongdoing in terms of the harms I have done to the community's valuable social practice. While this gives the duty to rescue a direction, it isn't the right direction.[90]

So, duties to rescue in the pond-like cases appear to have three features that are also characteristic of deontological constraints: priority, significance, and direction.

---

saying the imperiled person has a "claim" upon me. In arguing for a somewhat restricted sense of claim-rights, Judith Jarvis Thomson offers an example in which one may have a moral obligation to do something on behalf of another person even though that other person has no claim or right to one's compliance with that obligation: two students are attending a lecture and the first student's only pencil breaks. Thomson thinks the second student, who has plenty of extra pencils, has a moral obligation to lend one out, even though the first student has no claim against the second to do so (1990: 117).

In pond-like cases in which there is more than one imperiled person to be rescued, and one cannot rescue them all, some may deny that anyone has a claim up me, or that I have an obligation towards anyone (McKinsey 1981: 311). Feinberg, on the other hand, argues that while there is in such a case no one *in particular* whom I am obligated to rescue (the duty lacks a "determinate beneficiary"), I would violate the rights of every imperiled person in the water if I refrained from acting. My wrongdoing, in such a case, *may* also have a victim, namely the person who would have been saved had I acted (assuming there is such a person) (1987: 143-48).

[89] Consequently, and counter-intuitively, Arthur Ripstein argues that "rescuing others in emergencies should be seen as a duty to contribute to just distributive schemes, rather than as a duty owed to particular persons," and "[t]he crucial feature of any such duty…is that it is not owed to the particular person. Instead it is owed to society in general. That is, the duty to rescue is best thought of as a part of a more general duty to sustain just institutions….[I]t is a non-relational duty…..Those who fail to contribute breach a duty to society as a whole, rather than to the particular person who is not rescued" (2000: 766 fn. 20, 774-75). It is unclear whether Ripstein takes these conclusions to be about one's moral duties, as opposed to legal duties. Cecile Fabre dissents (2006: section 2.2).

[90] It might be noticed that the norms of the assumed practice of mutual aid themselves contain the directional aspect we seek to explain. It is, however, unclear why duly respecting the value of the practice requires one to heed the very norms that constitute the practice. Just as one can duly respect the value of baseball without playing baseball, one might be able to duly respect the value of the assumed practice of mutual aid without engaging in the practice itself. Or, even if one is engaged in the game of baseball, competing values may give one sufficient reason to violate its rules.

While these *deontic features* are plausibly interrelated—the priority requires a corresponding significance, and the direction supports that significance—it isn't clear how to explain any of them. Although special duties often exhibit such features, it is difficult to see how the relevant rescue duties can be understood to be special duties. So, a large gap separates what we have—the Millerian value—from what we want—a duty with deontic priority, significance, and direction.

Although these difficulties may seem to suggest that we should look elsewhere than the Millerian value for a differentiating feature, I think this would be a mistake. First, it is far from clear where else we should look. Whichever differentiating feature we choose, we will encounter the challenge of trying to explain why I owe it to the imperiled person in Shallow Pond to perform the rescue, but not to the person whom I might rescue through Heroes for Hamiltons. We must resist the temptation to beg the question, appealing to the idea that the imperiled person in the Shallow Pond has a claim to my rescue, or that my refusal to perform this rescue would be particularly disrespectful. Our task is to account for this person's claim or for this violation of respect. We may want to explain the imperiled person's claim upon me by pointing to what is at stake for that person in the Shallow Pond. The problem here is that it is difficult to see how any such personal value can even qualify as a differentiating feature. What, after all, is at stake for the imperiled persons in the pond-like cases which isn't also at stake for some imperiled person in a famine-like case? If, on the other hand, we appeal to some impersonal value, then we are confronted by the very challenges we face when appealing to the Millerian value: there remains a gap between the impersonal value and duties with the deontic features.

The deontic features of duties of rescue in the pond-like cases are therefore a general challenge for the moderate, one that isn't generated by the idea that such duties are special duties, nor by the idea that we should appeal to the Millerian value. In short, the moderate is in no position to evade the challenge of explaining the deontic features of the rescue duty.

To meet this challenge we must think about how features characteristic of deontological constraints and special duties can come to be attached to a duty to rescue a stranger for whom one has no particular, voluntarily assumed, responsibilities. I will first try to show how the Millerian value, which is the value of a social practice, can generate, or translate into, individual duties to perform rescues. I will then try to justify the significance of duties so understood, and to explain their directional aspect. Lastly, I will address the issue of deontic priority. Insofar as the object of the Millerian value is a social practice that doesn't extend to the famine-like cases, my resulting account will contain the needed differentiating feature.

3. From the Millerian Value to a Stringent Duty

Some have suggested that individual responsibilities for helping people in need derive from fair allocations of our collective responsibility for meeting needs generally. On such accounts, the importance of people's welfare generates a collective responsibility for promoting welfare, and the task is to show how this collective responsibility is fairly distributed among individuals. However, if the fundamental value is simply that of promoting welfare, or of meeting needs, it is difficult to explain how the

only appropriate allocation of responsibility requires such a prioritization of the pond-like cases. As in the augmented case of Bob's Bugatti, such a prioritization can get in the way of directly respecting the underlying value. We can avoid this general difficulty if the value at stake is not the general one of meeting basic needs, but is rather the much more specific value of the assumed practice of mutual aid. The Millerian value already provides this advantage, but we still need to find a plausible way to derive individual requirements from the value of the societal practice.

3.1 Unacceptability judgments

Miller solicits our appreciation of the practice's value by having us consider its absence. A society that doesn't respect the assumed practice of mutual aid will, I suspect, strike many of us as unacceptable. To say that something is *unacceptable* is to say more than that it is simply bad. It is bad for us that the laws of physics preclude cold fusion. This, however, is not unacceptable. It is also merely bad that there isn't an inexpensive pill that makes us, without side effect, happy and reasonable. A bad thing is unacceptable only if it is also in some sense unnecessary, preventable, or avoidable.[91] The absence of the assumed practice of mutual aid is unacceptable in part because this absence would be so unnecessary in the relevant sense. That absence would be the consequence of human choices, and if people only behaved differently—and they easily could—the good social practice would be realized.

[91] Here I state merely one necessary condition of unacceptability because I want to highlight the connection between the unacceptable and the avoidable. Spelling out all the conditions would take us farther afield. For a bad condition to be unacceptable in the normatively significance sense that is relevant, it may also be that the badness itself has to be somewhat egregious. An innocuous condition may be bad relative to what it otherwise might feasibly be, and yet still not be, in the relevant sense, unacceptable.

This connection between the value of the assumed practice of mutual aid and our unacceptability judgments is noteworthy. There are a great many things that are valuable or good that do not generate moral requirements. Moral requirements seem to concern practical standards of minimal decency, requiring us, in our actions, not to fall below that line. If, as I have suggested in the previous chapter, the standards of minimal decency are generally sensitive to costs of compliance, then our unacceptability judgments are in the same ballpark as our judgments about moral duties. Costs of compliance will affect the degree to which any bad condition we bring about is feasibly avoidable, and this will in turn determine whether the bad condition is unacceptable. Miller's observation, which is essentially a judgment about the unacceptability of specific conditions brought about by choices, is therefore a promising place to anchor an account of rescue duties. I expect that these connections will become clearer as we proceed.[92]

3.2 To practical norms applying to individuals

The first gap to address is that which separates unacceptability judgments from practical norms applying to individuals. Unacceptability judgments concern the kinds of things that both hold value/disvalue (whether instrumental or non-instrumental) and which are dependent upon choices (although not necessarily upon any individual choice).

[92] In relating moral standards to unacceptability judgments, I am guided in part by Scanlon's analysis of rights claims: "To claim that something is a right, then, is to claim that some limit or requirement on policy decisions is *necessary* if unacceptable results are to be avoided, and that this particular limit or requirement is a *feasible* one, that is, that its acceptance provides adequate protection against such results and does so at a tolerable cost to other interests….What rights there are in a given social setting at a given time depends on which judgments of necessity and feasibility are true at that place and time" (Scanlon 2003a: 99). Similarly, in "Rights, Goals, and Fairness," Scanlon writes, that "the case for most familiar rights—freedom of expression, due process, religious toleration—seems to be more concerned with the avoidance of particular bad consequences" in a way that is sensitive to "considerations of cost" (2003a: 34-38). Scanlon is usually addressing legally enforced rights, and so the consequences of a given "assignment of rights" might be assessed in terms of the effects of coercive enforcement. As he notes, his "own account emphasizes the value attached to rights for the sake of what they may bring rather than their value as signs of respect" (38-39).

So understood, the object of an unacceptability judgment may be a state of affairs, or it may be some specific feature of the world, such as a particular relationship or a practice. So as not to restrict possibilities, I will simply speak of the objects of unacceptability judgments as "conditions." What we are looking for is some way to derive individual practical norms from judgments about unacceptable conditions.

In considering how this might be done, we can think of ways in which criticisms of actions appeal to unacceptable conditions. There are, on the one hand, criticisms of an action that point out how the action itself causes the unacceptable outcome: throwing a lighted match into a bed of pine needles during a heat wave in Yellowstone is bad because it will cause a needlessly devastating fire. This, however, is not the relation we find between one failure to perform a pond-like case rescue and the unacceptable absence of the valuable social practice. The one failure to perform the rescue cannot be rightly criticized for obliterating the valuable social practice. The one failure to perform the rescue might not even weaken the relevant social practice, or, if it does weaken or undermine the social practice, this effect may be trivial.[93] Finally, when we consider the particular choice in my augmented case of Bob's Bugatti, it is difficult to see how, if I should let the one child perish, the overall outcome is even bad.[94]

In other cases, we appeal to unacceptable conditions in order to criticize an individual action that neither threatens to produce the unacceptable outcome, nor is even, on its own, particularly bad. These cases merit attention. I will offer an example of such

[93] Circumstances might, for example, allow me to keep secret my refusal to perform the rescue in Shallow Pond. Or, should my secret be discovered, I might still fortify the public sentiments supporting the social practice by confessing my refusal with a repulsive show of arrogance, or, oppositely, with groveling humiliation.

[94] If we do insist that letting the one child die in such a circumstance really is unacceptable, it looks as if we are just responding directly to the perceived wrongness of the particular action. Unacceptability judgments of this sort cannot help to explain the wrongness of the action.

a case that is very different from the pond-like cases and the deontic rescue requirements. This case is instructive because of a few relevant features.

Suppose that the following is a fact: as a direct result of people's careless disposal of household batteries—tossing them out with the trash—dangerous levels of toxic metals have seeped into the water table, a vital source of our drinking water. Filtering out these toxic metals is impractical. The public has known about these consequences for some time and, during this time, free means of safe battery disposal have been available at conveniently located recycling centers. Importantly, the dangerous levels of toxicity are the collective consequence of the negligent behavior of hundreds of thousands of residents. The effect of tossing one dead battery out with the garbage is trivial. Even the effect of tossing out a household's yearlong accumulation of dead batteries is trivial. For these individual acts, the resultant increase in the groundwater toxicity poses no significant harm to anyone. Such individual acts don't, on their own, even pose any risk of harm.

Still, of the person who carelessly tosses out batteries in this scenario, I think we want to say that some kind of criticism is both weighty and appropriate. One colloquial form that such a criticism might take is this: "It is because of people like you that the groundwater is dangerously toxic." I expect that such a criticism will deliver a forceful sting to anyone distressed by the needless toxicity of the groundwater.

The criticism takes a form that is quite common and, I think, readily intelligible. If we were to hear any of the following statements uttered among strangers we could readily appreciate the kind of accusation being leveled: "It is because of people like you that everyone has to lock their doors," "that the grades at this college are so inflated,"

"that there are still outbreaks of whooping cough," "that inhumane methods of factory farming are so prevalent," or "that those idiots might win the election." When the outcome appealed to is instead positive, we can similarly recognize the praise: "It is because of people like you that humanity still has a fighting chance," or "that Oxfam is now able to work in over ninety countries," etc.

Statements of the form "It is because of people like you that X," where *X* is some valued or disvalued condition, explicitly assert only a causal judgment, or a judgment that X is brought about by "people like you." The normative content is implicit. When statements of this form are intended as criticism, as opposed to praise, they contain the implicit normative judgment that the condition in question is unacceptable.

Such statements of criticism also, I think, contain the implicit normative judgment that the person in question is blameworthy for being among the people who are bringing about the unacceptable condition. If the person is blameworthy it is because of that person's attitudes or actions. It is people who act in certain ways who are causally responsible for bringing about the unacceptable condition. When people act in those ways without excuse, they are also blameworthy. So, when the criticism states, "It is people like you," instead of "It is actions like that," this implicitly identifies the problematic type of action, and it additionally implies that the person in question has no good excuse for performing an action of that type. The crucial point is that the normative content here is practical and personal, concerning the sorts of things for which individuals have personal control.

We notice, then, that the colloquial judgment "It is because people like you that X" implies a personal, practical norm which is somehow derived from the unacceptability

of some condition. This implication apparently holds, moreover, despite the full recognition that the individual action, for which the person is blameworthy, doesn't of itself bring about the unacceptable outcome. What we are looking for, then, is the principle that allows us to derive the personal, practical norm from the unacceptability judgment concerning X. We are looking for the *bridge principle* that connects unacceptable conditions to personal, practical norms.

When we express criticisms of this form, I suspect that we typically have already identified some type of action that, in aggregate, brings about the outcome in question. Moreover, part of what goes into making the outcome unacceptable is that the individual actions of the identified action types are, or were, so unnecessary. In the example given, the identified type of action is something like this: the careless disposal of household batteries. So, the criticism presupposes that, collectively, the careless disposal of household batteries is causally responsible for the pollution of the groundwater. The personal, practical norm is simply the norm against performing individual actions of that type. If you do perform such actions, and you lack a good excuse for doing so, then you are blameworthy, and the criticism in question applies: it is because of people like you that groundwater is so polluted. Generally, then, the suggested bridge principle is simply this: don't perform individual actions of the identified type that collectively bring about the unacceptable condition.

The unacceptable condition that Miller envisions differs in several potentially important respects from the unacceptable condition in the battery disposal case. One is that, in Miller's observation, the unacceptable condition is unrealized. There currently exists a social practice of rescuing one another in pond-like cases. This difference may

not be deep. We might suppose, for example, that despite current levels of toxicity, many people continue to toss out their batteries with the trash, and, should this trend continue, the dangers to public health will advance from bad to catastrophic. In such a case, the full catastrophe isn't realized. This merely potential catastrophe might still figure into a criticism of a similar form. Of the person who carelessly disposes batteries, it might be said: "It is because of people like you that the groundwater pollution *threatens* to become catastrophic" (similarly: "It is because of people like you that those idiots might win the election"). Alternatively, if conscientious behavior is widespread, another criticism can be leveled: "It is only because others don't allow themselves to act like you that the toxicity isn't catastrophic."[95]

---

[95] In this last criticism, the implicit causal judgment concerns an unrealized outcome; its realization may not even be a live threat. When the unacceptable outcome isn't realized, the causal judgment identifies a class of actions that needlessly *would* bring about the unacceptable outcome. The interpretation of "would" is tricky.

We might characterize a class of actions in terms of principles of action applying in specific circumstances, where such principles identify behavior to be performed in a specified circumstance. We can say that such actions *would bring about* whatever condition is created simply by everyone's full compliance with the given principle, with everything else remaining as otherwise expected. We can say that the *costs of compliance* are the costs for individuals, as individuals, in whatever circumstances the principles apply, encompassing various levels of others' compliance. (Notice, then, that the levels of compliance we should assume when assessing compliance costs differ from the compliance levels we should assume when determining what conditions "would be brought about." It is a mistake to think that we must latch onto the same level of compliance to answer both questions.) Whether the outcome of such actions is *unacceptable* depends (at least) on these costs of compliance and the value/disvalue of the condition that such actions would bring about, where both the costs of compliance and the value/disvalue of the outcomes is assessed relative to those of alternative principle-specified actions. To illustrate, consider the principle of action under which a person donates $10 to Heroes for Hamiltons whenever this would rescue another person. In the imagined world of Distant Rescue, the costs of compliance to individuals are very high because such a principle applies even at very low levels of compliance. The condition that the actions would bring about would be good (Heroes for Hamiltons would rescue everyone who can be rescued for $10), but it would not be any more valuable relative to the following, slightly less specific, alternative: rescue a person whenever this can be done at a very low cost to oneself. If, moreover, Heroes for Hamiltons would have sufficient funds to rescue everyone it identifies if everyone gave just ten percent of their income to humanitarian projects, the outcomes associated with a far less onerous principle would be just as favorable.

I suggest that the assessment of unacceptability depends *at least* on such considerations because there are undoubtedly other relevant considerations. There may also be considerations of fairness in how the principle of action distributes costs. There are also considerations of costs and value associated with less-than-full compliance levels (e.g., while full compliance may be very valuable, one's own compliance under conditions of partial compliance may be futile or even deleterious). Similarly, sometimes the

This last form of judgment suggests further differences when the unacceptable outcome isn't realized. First, it may be the conscientious sacrifices of others that are preventing the unacceptable outcome. Second, others may be drawing benefits from the absence of the unacceptable outcome. So, as an existing good, the absence of the unacceptable outcome may be something that others are presently relying upon. If, for example, the groundwater is not yet polluted, it is likely to be a good that people depend upon, which, if anything, only increases the significance of polluting it (Goodin 1985: 124-25). Third, if you are among those drawing a benefit from the unacceptable outcome's absence, and this is sustained by the good behavior of others, then your bad behavior might also violate some constraint against freeriding.[96]

These considerations offer further normative resources either for guiding the unacceptability judgment, or for deriving the practical norms against what is essentially the same kind of behavior, namely the actions identified as bringing about the unacceptable outcome. In the imagined case, these additional considerations either support the conclusion that the yet-to-be-realized catastrophic levels of pollution are unacceptable, or to they add force to the criticism of those actions that threaten to bring about that unacceptable outcome.

---

value/disvalue of acting in a given way may be sensitive to people's expectations of how people will behave. If so, then we shouldn't assume the particular value/disvalue of acting in a given way in the actual world is accurately represented in a world in which expectations are very different. In our assessments of unacceptability, we may therefore need to fix expectation levels somewhat artificially so that they do not vary with the differing compliance levels we must imagine.

[96] Some may view a constraint against freeriding as central, fundamental, or at least independent from the kind of proposal I am offering (as in Rawls's principle of fairness (1999a: 96-7)). I regard the constraint against freeriding as a supplement to the general proposal I am offering, and one that is perhaps even derivable from the general proposal. When it is criticizable, we can regard freeriding as a more general type of action that brings about, or threatens to bring about, unacceptable outcomes. An action may belong to more than one type, and each type may be criticized in terms of productive relations to different unacceptable outcomes. Careless battery disposal might be criticized as the kind of action that threatens to bring about catastrophic water pollution, and, more generally, as the sort of action that frustrates, via freeriding, the realization of public goods.

We needn't regard these criticisms, or the norms they express, as moral criticisms. It suffices here if these criticisms are common, intelligible, and legitimate; that they express disapproval for the violation of legitimate norms that assign responsibilities to individuals; and that the disapproval on such grounds is recognizably significant. In the battery disposal case, if I throw out my batteries with the trash, others may rightly say that it is because of people like me that the groundwater is getting contaminated. This criticism plausibly expresses both a reason for others to blame and resent me, and a reason for me to change my behavior. The criticism therefore carries both interpersonal significance, and significance for me as an individual deliberating about what to do.

3.3 *Adding deontic direction*

In this example, the value at stake (the potable groundwater) is a public good which everyone in the community has reason to respect. The criticism towards me for my careless battery disposal expresses a norm that allocates to each individual in the community a specific responsibility with regard to that public good. However, although that responsibility is personal in that it applies to me as an individual, it registers a non-personal, or impersonal, complaint in the specific sense that its force is not captured by any complaint that is properly made on any particular individual's behalf.[97] In this sense, then, the norm that applies to me lacks the kind of directional aspect associated with

[97] Here it may be said that the complaint is personal because it is made on *each person's* behalf. I take it, however, that the complaint is more significant for that reason, that is, for the reason that it is made on many people's behalf. In this sense, the significance of the complaint isn't represented by any complaint that can be made on behalf of any particular person. So, as the complaint of the aggregate, the special significance of the complaint has an impersonal aspect. For my interpretation of the personal/impersonal distinction, where the personal claims are those that can be made on an individual's behalf, I believe I follow Scanlon in what Derek Parfit has called his "individualist restriction" (Scanlon 2003b: 429). It is important to me to try to show how such impersonal values can come to support personal claims and the particular moral significance that personal claims in particular seem to carry.

prototypical deontological duties. While proper disposal may be something that I owe to my community, it isn't something that I owe to any individual in particular.

This impersonal complaint might, however, come to acquire a further personal aspect. Should a particular child become ill on account of the contaminated groundwater, my irresponsibility might then be said to have a victim. This particular child might then be said to have a personal grievance against me. Here however it might be noted that, while the child has a personal grievance against me, her grievance isn't legitimately directed towards me *in particular*. Rather, her grievance is directed towards all those who, like me, have negligently contributed towards the groundwater pollution.

With a small adjustment the directional aspect can be strengthened so that the victim's grievance is against me in particular. If I toss a bag of garbage into the river, then it might be said that it is because of people like me that our waterways are so polluted with trash. If, however, a year later my rancid bag of garbage washes up and befouls someone's newly acquired riverfront property, then my irresponsibility comes to have that person as a particular victim. She has a complaint to raise on her own behalf against me. It is then, I think, sensible to say that I (in particular) have wronged her (in particular).

So, the unacceptable outcome often has particular victims. When it is on account of the interests of such victims that the outcome is judged to be unacceptable—and responsibilities for its avoidance are therefore allocated—then such people have what we might call a non-incidental interest in the avoidance of the unacceptable outcome. The bridge principles in effect assign to individuals portions of responsibility for protecting

these people's interest in the avoidance of the unacceptable outcome. Towards these people we say the fulfillment of the individual responsibilities is owed.

As an analogy, consider a municipal law assigning to each homeowner the responsibility to clear the snow from the public sidewalk in front of his or her house. Among the many ways to deal with the snow, the city has chosen this one. Now I, as an individual homeowner, have a responsibility. Facing that responsibility, I might ask, "Why should I comply?" One reason is this: otherwise the city might fine me $25. $25, however, is just a small amount, and I might rather pay the fees than perform the task. There is, however, further reason to shovel the snow. Suppose that compliance levels with this ordinance are low. This means that the sidewalks are a mess and, as a result, every winter several people slip and break bones. In this case, another consideration might grab my attention: if I fail to clear my share of the sidewalk, it will be because of people like me that the sidewalks are so unnecessarily dangerous, and that, from time to time, pedestrians are seriously inconvenienced and injured through falls. Alternatively, suppose the compliance with the law is high. As the result, I can stroll with ease midwinter along the ice-free pavement. Now, even if I don't clear my portion of the sidewalk, the sidewalks are still just as nearly ice-free. However, I might rightly be accused of freeriding.

Both kinds of suggested complaints are, in a sense used earlier, impersonal. That is, the force of the complaints that I am freeriding on a public good, or that it is because of people like me that the sidewalks are so unnecessarily dangerous, is not captured by any complaint that can be made on behalf of any particular individual. They are made, perhaps, on behalf of the community as a whole, of everyone in the community, or of

everyone who has a legitimate interest in the public good, or who has done his or her share of the collective responsibility. Though complaints that are in this way impersonal lack the sort of character and significance I associate with prototypical deontic requirements, they may still be significant. They might even be said to address moral wrongdoings or moral failings.

To add the peculiar directional aspect of a deontic duty, suppose that a failure to scrape my portion of the sidewalk inconveniences a disabled neighbor up the street. On account of the ice, the disabled man finds it difficult to operate his wheelchair. Under such circumstances, my reason to fulfill my share of the snow shoveling might come to have deontic direction. When I now consider why I should shovel the snow, the potential complaint of the disabled man now grabs my attention. Shoveling the snow is something that I owe to him.

I intend this analogy only to provide an illustration. It is not a demonstration since there are alternative explanations for why the situation of the disabled man changes the character of my responsibilities. The analogy is meant to suggest a way in which an initial allocation of responsibilities, parallel to the kind generated by unacceptability judgments and the bridge principle, is compatible with deontic direction, a feature that may come to take the foreground when we then reflect upon the responsibilities.

The deontic direction does not arise from thin air. In the snow-shoveling case, concern for the disabled man's interests is at least part of what calls for an allocation of responsibility in the first place. The disabled man is just a more pressing case of the interests pedestrians generally have in accessible pathways. These interests, alongside the feasibility of collectively securing them, yield a collective responsibility; they ground

an entitlement for some relevant provision. This entitlement gives individuals, and the disabled man in particular, claims upon others for these provisions. When this responsibility is legitimately allocated (in this case through the municipal law), giving to individuals specific responsibilities, then the disabled man's claim becomes more specific: it becomes a claim upon those particular individuals to meet their specific responsibilities. So, the directional element originally connected the disabled man and his community, but it gave him a claim towards no one in particular. The allocation of that responsibility then connected the disabled man's claim to the particular individuals to whom the specific responsibilities were assigned. So, when I fail to clear the snow off my portion of the sidewalk, the disabled man can rightly identify me as one who is to blame for his failing to receive the provision to which he is entitled. His disability only increases the significance of his loss, and therefore the seriousness of my failure.

This analogy also suggests the sense in which the disabled man has a *non-incidental* stake in my fulfillment of my snow-shoveling responsibilities. Interests like those of the disabled man in an accessible pathway are, at least in part, and perhaps in the aggregate, what call for the allocation of snow-clearing responsibilities in the first place. This stands in contrast to the incidental stake that a local seller of snow shovels has in my compliance with the ordinance. Though the merchant is, in her sales, adversely affected by low compliance, she is not therefore wronged, for it is not for the sake of the merchant's sales that the community has a responsibility to deal with the snow.[98]

[98] The general issue here is "the problem of third party beneficiaries" that faces interest-based theories of rights (Hart 1955: 180). I am suggesting that the interests that generate directed duties are those kinds of interests for which a further condition holds: the failure to make provisions for such interests generally would be, perhaps because of certain impersonal values at stake, unacceptable. In its appeal to a more general, impersonal value that lends moral significance to the interests of individuals, my proposal is similar to Joseph Raz's view that a claim-right is "justified by the service it does to the interest of the right-holder" even though the importance of the interest—including, I take it, the sufficiency of that interest to

Returning to the pond-like cases, we can see that Miller's observation can provide an analogous basis for individual rescue duties with deontic direction. Miller's observation solicits a judgment about the unacceptability of the absence of the assumed practice of mutual aid. Since the practice just is the general respect for the norms requiring one to perform the rescue in the pond-like cases,[99] it is easier to identify the unnecessary action type that is, or would be, causally responsible for an unacceptable state of affairs. Here, the behavior to be criticized just is the failure to respect the norms of the practice in question. The plausible bridge principle that links unacceptability judgments to practical norms simply prohibits such failures, thereby assigning to everyone the individual responsibility of performing the rescues in the pond-like cases. Since, however, there will always be a particular person who has a significant, non-incidental stake in one's fulfillment of that responsibility—namely the imperiled person—such a duty takes on a directional aspect: it is owed to whomever it is who happens to be imperiled. If I should find myself in the position to perform the rescue in a pond-like case, the imperiled person plausibly has a claim upon me to fulfill my responsibility. Such a person's stake is non-incidental in the same sense that the disabled man's stake in my snow shoveling is non-incidental. As with the disabled man's claim, it is the claim of the imperiled person that jumps to the foreground when I consider why I should perform the rescue.

---

ground rights—may itself derive from more general importance (to the interests of others, or to the public/collective good) of providing for that kind of interest (1988: 181-82, ch. 10).

[99] Characterizing a practice in this way distinguishes the relevant value both from the value of an assignment of rights and the consequences thereof (as in Scanlon's "Rights, Goals, and Fairness" proposal (footnote 92 above)), and from the value of morality itself assigning a certain "status" to individuals (as proposed by Kamm in 2001: ch. 10). The former proposal seems to be a better fit for legal rights, where an assignment of such rights will include some kind of enforcement. The latter proposal strikes me as too thin, or at least as not capturing the kind of loss to the community that Miller identifies.

With this deontic direction, the duty to do one's part in blocking the unacceptable outcome takes on a distinctively moral character.[100] This deontic direction adds to one's responsibility the feature that it is something that one owes to a particular person. Failing the responsibility now has a particular victim, an individual whom one will have wronged. Once our responsibility for doing something is owed to someone, a failure to fulfill that responsibility disrespects the particular person, depriving the person of something that we owe to him or her. For the morally conscientious person, such considerations are significant. The morally conscientious person is particularly concerned not to be the one who deprives another person of that to which he or she is entitled. This is all the more pressing when, as in the pond-like cases, that to which the person is entitled is of vital importance to him or her.

3.4 *Explaining deontic priority*

Deontic direction may account for at least some of the priority of the responsibilities that have deontic direction. When my reason for not polluting, or for shoveling the snow, takes on deontic direction, then my failure to respect the reason constitutes a failure to respect the person to whom my responsibility is now connected. If, then, I have further and weighty reason simply to respect other people; or to render to others what I owe to them; or to maintain my relations of "mutual recognition" with them, acting only in ways that are justifiable to them; then this deontic direction makes my polluting, or my failure to shovel the snow, more consequential. Deontic direction

[100] Deontic direction puts the responsibility within the range of Scanlon's "morality in the narrow sense" (1998). J. J. Thomson writes, "There are endless things that we ought to do though others have no claim against us that we do them….The realm of rights is squarely within the morality of action, but it is not identical with it" (1990: 117), cited in Sreenivasan 2010: 472. R. Jay Wallace argues that what I call deontic direction is the characteristic feature of moral obligation (2012, unpublished).

thereby ratchets up the significance of my responsibility. This added significance may increase the priority that I should give to snow shoveling over the other things I have reason to do.

This sense of added priority, however, does not independently account for the feature of deontic priority, or at least not in its full-fledged, "paradox"-generating sense. While deontic direction is an important part of the moral significance of our reason to perform the rescue in the pond-like cases, this deontic direction doesn't account for the way in which such duties still apply in scenarios like my augmented case of Bob's Bugatti. It may well be true that *if* I owe it to one imperiled person to rescue her, I cannot rightly opt to do anything else instead. However, all I have yet provided is an explanation why I owe it to an imperiled person in Shallow Pond to rescue her, and how such an obligation might derive from the value of a more general practice, like that of mutual aid in the pond-like cases. I still need to say why I might owe it to the imperiled person to rescue her even in the special sorts of circumstances here imagined, where I may also have opportunity to better serve many other equally imperiled people as an alternative. In other words, while we might concede that, in many pond-like cases, there is a special duty, owed to the imperiled person, one needn't concede the full-fledged deontic priority of such special duties. Alternatives that would allow me to do so much better for so many others far away might seem like they should defeat any duty to the person nearby.

If I am to make the case for the full-fledged deontic priority of the duty to rescue in the pond-like cases, then I must argue that I owe it to the nearby imperiled person to rescue her *in the face of* many other things that I might alternatively do that would,

perhaps from some impersonal point of view, have equally good if not better consequences. Since I seek to maintain a parallel to special duties, and at least some special duties plausibly do have this kind of deontic priority, such an argument is worth an attempt.

On my proposed account, the deontic priority of a rescue duty will derive from the more general value of the practice that calls for such a rescue. The assumed practice of mutual aid calls for people to perform the rescues in the pond-like cases. When I suggested that the absence of this practice would be unacceptable, I had in mind the condition in which individuals acknowledge no norms requiring them to help the imperiled person in cases very much like Shallow Pond. This is the sort of unacceptable condition that Miller characterized. What seemed unacceptable was the condition in which a person violates no recognized norm in walking past the child drowning in a shallow pond.

To say that such rescue duties also have the kind of deontic priority that generates paradoxes of deontology, we would need in addition to say that changes to the assumed practice of mutual aid that were more modest and reasonable are also unacceptable (supposing that the assumed practice does include rescue duties of a "paradox"-generating kind). We might imagine that, while there is a practice of performing the rescue in Shallow Pond, there is no such recognized norm of rescuing the child when, by letting the child perish, one might more effectively help imperiled people overseas. If, however, such a deviation from the assumed practice of mutual aid is still unacceptable,

then we might appeal to this further judgment to defend the deontic priority of the rescue duty.[101]

When we ask these questions, I don't assume that full-fledged deontic priority of the rescue-duty will always be vindicated. Some changes, however, do strike me as unacceptable deviations from the assumed practice of mutual aid. For example, I find unacceptable the practice that licenses the refusal to perform the rescue in Shallow Pond so long as one promises to give an extra twenty dollars to Heroes for Hamiltons. Such a change would, it seems to me, greatly diminish the valuable sense of community we have with those around us, and we would lose this value gratuitously. There is a sense in which the integrity of the practice would be deeply damaged if it allowed for such potentially pervasive exceptions. Highly unfortunate global circumstances, such as those that allow Heroes for Hamiltons to successfully operate, would, through such exceptions, undermine the assumed practice of mutual aid (one could always, and very easily, opt out of the pond-like rescue), and the loss would not be offset by any general gain for imperiled people far away (since pond-like cases are so rare, the upshot of people opting out of pond-like case rescues won't be any great increase in the donations on behalf of the distant needy). It is also not as if such exceptions would greatly improve our

[101] It might be thought that a similar strategy could be used to explain the deontic direction of rescue duties in pond-like cases. That is, since the norms of the assumed practice of mutual aid themselves contains the very directional elements we seek to explain, we might appeal to the unacceptability of a slightly altered version of that practice in which no such directional components are recognized. While I would share the unacceptability judgment, I suspect that the absence of the directional components are unacceptable because of the significant, non-incidental stakes that the would-be beneficiaries have in the fulfillment of the allocated responsibilities to perform the rescues. In other words, in this case the appeal to the unacceptability judgment would miss the deeper explanation of why the absences in question are unacceptable. Generally, I am happy to accept a further explanation of an unacceptability judgment if a plausible and adequate one is available.

relationships with those far away; nor would greatly increase their prospects of being rescued by us.[102]

Things are less clear to me, however, when I think about my augmented case of Bob's Bugatti. Here I can see that a much stronger case could be made for allowing an exception to the duty of nearby rescue. If, however, we allow exceptions for some extreme cases, or if we just conclude that one's obligations in such cases are unclear, then we don't sacrifice the parallel to special duties and deontic requirements generally. Such exceptions would be like a catastrophe exception of a kind that plausibly belongs in many moral principles. I am required not to tell a lie—*unless* doing so is necessary to prevent some catastrophe (perhaps telling an otherwise inconsequential lie to a passerby is the only way to prevent someone from losing a leg). In many apparent pond-like cases there is room for disagreement, and the success of my proposal doesn't depend on its ability to resolve each of the hard cases.

Let me review. I am following Miller in supposing that what is peculiar about the pond-like cases is the value of the practice in which people, even complete strangers, are prepared to help one another in such cases. Miller's observation looks promising because it identifies a third value that might help explain the stringency of the pond-like cases, and do so in way that fits within a non-voluntarist account of special duties. The challenge is to explain how the value of such a practice can generate an individual duty with the deontic features. To meet this challenge, we can look to the practice's absence

[102] A parallel might be drawn to other special duties exhibiting some degree of deontic priority. Suppose that I am permitted to break my promise to you so long as this ensures that two other agents are able to keep their comparably weighty promises to yet other people. Such an exception seems somehow to threaten the integrity of promising, along with the valuable kind of guarantee that promising can (at least as ordinarily understood, where it doesn't allow such exceptions) establish between you and me. At the same time, the loss that such an exception would introduce isn't replaced by any sufficiently important gain. It isn't plausible that such a change to the assumed practice would greatly improve either the value of that practice, our relations with one other, or any other important value that is instrumentally gained.

as a condition that is not only bad but also unacceptable. Unacceptable conditions are commonly taken to license criticisms against the actions that bring about, threaten to bring about, or would bring about, such conditions. In the right circumstances, the standards expressed by such criticism contain deontic direction, ratcheting up the significance of these standards and lending to them a distinctively moral character. Finally, I have suggested how to explain deontic priority, or at least some degree of it.

4. The Pond Question Revisited

When I consider the plausibility of my proposal, I am nagged by the feeling that something is amiss with my proposed explanation for why I must rescue the drowning child in front of me. There is a range of appropriate responses to the question "Why must I rescue this drowning child?" but my proposed answer—which appeals to the value of a social practice, or the unacceptability of its absence—does not seem to be the kind of answer one should give. In the previous chapter, I argued that this apparent awkwardness is a bullet that the moderate will just have to bite.[103] Here I want to say why that awkwardness should not deeply trouble us.

The appeal to the value of the assumed practice of mutual aid may seem strange in part because alternative answers seem more fitting. These alternative answers, however, may seem more fitting because, in most ordinary contexts, they would be more fitting. They would be more fitting on account of conversational norms of the sort that Paul Grice has discussed (1989: ch. 2). However, we should only be troubled by this appearance of alternative fitting answers in two cases: first, if there seems to be an

[103] Ch. 4, sect. 3.3, pages 112 ff.

alternative answer that is more fitting in the context of the philosophical question we are addressing, or second, if my proposed answer is somehow incompatible with the fittingness of these other answers for those other contexts in which they truly are fitting.

It isn't clear that either of these cases obtain. The specific context in which I am providing an answer to why I must rescue the drowning child is that of the differentiation question, which asks for an explanation of why I must rescue that child, but not the child who might perish for want of a similarly small sacrifice to Heroes for Hamiltons. In this context, most of the explanations that would be fitting for other contexts in which I am asked to explain why I must rescue a drowning child simply don't suffice. Many such answers fail to identify a feature that differentiates the Shallow Pond from Distant Rescue. Other answers are unsatisfying because they seem ad hoc, failing generally to differentiate pond-like cases from famine-like cases. Some answers are unsatisfying because they point to a difference, like physical proximity, for which it is hard to understand why it should possess such moral significance.

My proposed answer, moreover, is compatible with many of these attractive answers of other contexts. Consider the answer that I should rescue the drowning child "*because she is right there in front to me*." Let us interpret this answer as pointing to the imperiled child's physical proximity to me. If such physical proximity is a characteristic feature of pond-like cases, we might regard this answer simply as expressing the appropriate sensitivity to the features that make a rescue scenario into a pond-like case (or, the features that make it so unacceptable for a community to lack the practice of mutual aid for the pond-like cases). My account attempts to justify the significance that we ascribe to the pond-like cases relative to the famine-like cases.

It may be that the morally conscientious deliberator is directly sensitive to general moral principles, principles that take for granted the sort of content provided by my proposal. It may be, for example, that the morally conscientious person is primarily concerned to avoid depriving another person of anything to which he or she is entitled. Facing a pond-like case, this deliberating agent simply takes for granted the imperiled person's entitlement to help. Such a consideration, as opposed to the deeper explanation of that consideration, is all that he or she need attend to. What my proposal then does is provide the deeper explanation of that salient consideration, explaining why entitlements are distributed as they are taken to be.

So, within a pond-like case, the morally-sensitive deliberating agent obviously does not need to recount consciously the philosophical justification of the moderate's response to the differentiating question. It is sufficient for the deliberating agent to be sensitive to the sorts of considerations that, in this world, reliably indicate that she is in a situation governed by the assumed practice of mutual aid, or that the imperiled person has a claim upon her to be rescued.[104]

We can add, finally, that the fittingness of these simple answers to the pond question is what we might expect if my philosophical justification of the moderate's position is correct. The Millerian value, which is at the heart of this justification, concerns the valuable sense of community one feels with others whom one might encounter. That valuable sense of community goes beyond the mere expectation that

[104] Further questions might be asked about what these reliable indicators are, and which specific factors prompt us to either attribute or deny that the imperiled person has a claim upon us to perform the rescue. Analyzing the descriptive factors that distinguish pond-like cases from famine-like cases is not a project that I pursue in this dissertation. This project is central to Peter Unger's augmentation of the famine-relief argument; it is also central in Frances Kamm's replies to Unger and in her proposals about the significance of distance.

those around me would rescue me or my loved ones from serious peril, especially if such could knowingly be done at no risk and at very little cost. Some of that sense of community would be lost if, in the pond-like scenario, my potential benefactor did not take the salient fact as sufficient to ground my claim to be rescued and to motivate action. The value that is lost is similar to the value lost in the relationship between parent and child if, before attending to his child's needs, the parent should always have to recount to himself the value of the institution of child-rearing.

5. The Differentiation Question Revisited

As with the pond question, there is also a remaining concern about how my proposal answers the differentiation question. Here I want to address that concern and to state more succinctly why it is, on my proposal, that I must rescue the imperiled person in Shallow Pond even when I am not required to rescue another person through Heroes for Hamiltons.

On my proposal, we might say that the crucial differentiating factor that separates pond-like cases from famine-like ones is the Millerian value, the value of the assumed practice of mutual aid. It is the assumed practice itself that differentiates the two sorts of cases, since it requires action in the pond-like cases without similarly requiring action in the famine-like cases. The value of that practice then gives rise to individual duties to respect its norms. This answer, however, will not be satisfying unless it is understood why an alternative valuable practice, whether actual or hypothetical, cannot similarly generate duties to rescue imperiled people in the famine-like cases.

This challenge has force insofar as it is easy to imagine practices more demanding than the assumed practice of mutual aid that, if realized, would be highly valuable. A more demanding practice could extend to all people the fellowship and solidarity that the assumed practice of mutual aid only facilitates for restricted circumstances. The more demanding practice would also add security to many people, and it would relieve certain of the inequalities sanctioned by the assumed practice. Moreover, while the value of the assumed practice of mutual aid generates duties on account of the unacceptability of its absence, we might also think that the absence of the more demanding practice is also unacceptable. After all, aren't the conditions that allow an organization like Heroes for Hamiltons to operate clearly unacceptable? Isn't there something outrageous about the circumstance in which I knowingly allow some other human being to perish for want of the ten dollars that I am about to spend on a fancy cocktail? Using the terms of the Millerian value, we can say that my purchase of the cocktail signals a striking lack of fellowship between me and fellow human beings.

To clarify this challenge, we need consider the *non-moderate rescue practice* according to which I am required to perform rescues in famine-like cases just as the assumed practice of mutual aid requires me to perform rescues in the pond-like cases. It is this particular alternative that directly threatens the moderate's resistance to the demanding conclusion.[105] The non-moderate practice is unrealized, and therefore merely hypothetical, insofar as its system of norms is not generally recognized and respected.

[105] It is open to the critic to try to describe valuable alternative hypothetical practices that conflict with the assumed practice of mutual aid. While this wouldn't directly challenge the moderate's denial of the demanding conclusion, it could still challenge the moderate's answer to the pond-question. As I see it, the most promising candidate here would be a slight variant of the assumed practice whose only difference would be that it didn't require rescue in extreme scenarios like my augmented case of Bob's Bugatti. I touched on this possibility above.

The moderate must explain why we are not required to respect the norms of the hypothetical non-moderate practice on account of the unacceptability of its absence.

The moderate's best response, as I see it, is to concede that the non-moderate rescue practice would be valuable, and to further concede the unacceptability of its absence. The moderate should insist, however, on interpreting this unacceptability in a certain way. Recall that unacceptability judgments imply the feasibility of avoiding that which is judged to be unacceptable. In the example cases I discussed, the unacceptability of a condition turned on the fact that the bad condition in question is, or was, so easily avoidable: the bad condition can be avoided in a way that doesn't saddle anyone with excessively great responsibilities. So likewise, if we judge the absence of the non-moderate practice to be unacceptable this is only because we think that the practice could be realized without excessively burdening anyone. We could, for example, be thinking that if only *many* people acted *a bit* more conscientiously an organization like Heroes for Hamiltons could not exist; the non-moderate practice would never overwhelm any one person with extremely high costs of compliance because no one would be perpetually confronted by opportunities to perform additional rescue at small costs.

While these concessions accept the plausible parallels between the assumed practice and the non-moderate practice, they do not yield any parallel duty to act on every rescue opportunity presented in the famine-like cases. The concessions rather suggest that we can use a similar form of reasoning to understand that we also have duties concerning imperiled people in famine-like cases (and the moderate should never have supposed that our responsibilities for the needs of strangers are exhausted by the norms comprising the assumed practice of mutual aid). The licensed criticism would apply to

those unwilling to make even modest sacrifices on behalf of imperiled strangers far away, for (on the assumptions described) it is because of people like this that we cannot feasibly realize the non-moderate rescue practice, with all benefits to fellowship and security against peril that such a practice would provide. The corresponding requirement would only be a duty to act in those ways that, if many others did the same, we could feasibly realize the non-moderate rescue practice. For many people, this might just mean a requirement to be a bit more conscientious concerning the welfare of people far away.

The differentiation question asks why I am required to rescue the person in Shallow Pond even when I am not required to rescue someone through Heroes for Hamiltons. On my proposal here is what the answer should be. Moral duties are limited by their sensitivity to the cumulative costs of compliance. In particular, my duties to rescue others are limited by the cumulative costs to me of performing those rescues. This limitation blocks any requirement upon me to keep rescuing people through Heroes for Hamiltons so long as I can continue to perform such rescues at individually small costs to myself. It doesn't, however, block a more limited requirement to perform the rescue in the pond-like cases, where these cases are distinguished from the famine-like cases insofar as only the former are governed by the norms of the assumed practice of mutual aid. Were I to fail to rescue the person in Shallow Pond, I would be violating my responsibility to avoid performing an action of a type which collectively brings about, or would bring about, an unacceptable outcome. The unacceptable outcome that would be brought about by refusals to rescue people in the pond-like scenarios is the absence of the assumed practice of mutual aid.

6. Rescue Duties as Special Duties

In this chapter I have focused on the challenge of explaining the deontic features of the duty to rescue in the pond-like cases. I observed that a similar challenge would be faced in attempting to account for special duties. The special duties I have towards my children, for example, are in some ways similar to the duty I have to the imperiled person in Shallow Pond. That imperiled person has a special claim upon me that isn't shared by others equally imperiled whom I might also rescue; my children have special claims upon me that aren't shared by others whom I might similarly benefit. My corresponding duties to my children, similarly, seem to feature the kind of priority, significance, and direction of my duties to rescue in Shallow Pond. The tempting suggestion, therefore, was that I might borrow from the accounts of special duties to understand the duties to rescue in the pond-like cases. I did not, however, find what I was looking for. Although the discussions of special duties I found most suited to my purposes where those of Samuel Scheffler, I ultimately did not find them to provide the insight I needed. I now want to address more fully Scheffler's proposal, as it seems to be, in these respects, the nearest competitor to my own. I will close by considering the extent to which it is helpful to view the stringent duty of rescue in the pond-like cases as a special duty, and, in particular, how we can close the loophole that allows the problem of over-demandingness to arise.

6.1 Responding to Scheffler

In Scheffler's terminology, special duties, as opposed to general ones, are "duties that we have only to those particular people with whom we have had certain significant sorts of interactions or to whom we stand in certain significant sorts of relations." Most relevant to my own project is the subset of such duties that Scheffler calls the "associative duties." These are the special duties which "members of significant social groups and the participants in close personal relationships are thought to have to one another" (2003: 50-51, 68). I am particularly interested in Scheffler's proposal for how such duties arise.

According to Scheffler, to value, in a non-instrumental way, my relationship with another person "just is, in part, to see that person as a source of special claims in virtue of the relationship between us" (1997: 196). Moreover, if we are "correct to [so] value our relationships—insofar as our relationships are valu*able*—they are indeed sources of reasons" (2010: 110). So, according to Scheffler, a non-instrumentally valuable relationship that I have with another person makes this other person a source of special claim upon me. The valuable relationship "transforms the needs and desires" of this other person into special reasons for me to act on behalf of him or her in suitable contexts (110).

Scheffler's account is most questionable when he tries to explain how these reasons to act on behalf of the other person become duties to so act. Scheffler draws out the challenge by considering the asymmetry between relationship-dependent reasons and project-dependent reasons. Drawing a parallel between these two types of reasons, Scheffler claims that "if one values a personal project non-instrumentally, then one will see oneself as having reason to devote special attention to that project." If, moreover,

that personal project really is valuable, then it presumably is a source of special reasons. The asymmetry is that we are often "morally required" or "obligated" to act on the relationship-dependent reasons, but we are not so required to act on the project-dependent ones (2004: 258-60, 2010: 108-09).

After rejecting a few proposal for explaining this asymmetry, Scheffler latches onto the idea that relationship-dependent reasons are importantly different because of their "shared" or "interlocking" characteristics. Here is how he describes the important characteristic:

> A valuable relationship transforms the needs and desires of the participants into reasons for each to act on behalf of the other in suitable contexts. At the same time, it gives each of them reasons to form certain normative expectations of the other, and to complain if these expectations are not met. In particular, it gives each of them reason to expect that the other will act on his or her behalf in suitable contexts. These two sets of reasons—reasons for action on the one hand and reasons to form normative expectations on the other—are two sides of the same coin. They are constitutively linked and jointly generated by the relationship between the participants. (2010: 110)

Scheffler seems to be suggesting that if *A*'s relationship with *B* gives *A* reason to act on *B*'s behalf, then *B* has reason to complain if *A* does not do so. The further implication, it seems, is that if *B*'s entitlement to so complain against *A* shows that *A* owes it to *B* to act on *B*'s behalf. If *A* so owes this to *B*, then the action is morally required; *A* wrongs *B* if she fails to perform the action.

Suppose, however, that I have a one-of-a-kind interpersonal relationship with my next-door neighbor, but that it is only slightly valuable. The slight value of this relationship may still generate some relationship-dependent reasons, but I expect that these will be weak, and rarely, if ever, decisive. So, then, if this slightly valuable relationship generates a slight reason for me to fire up the grill and invite him over from

across the driveway, it is doubtful whether any corresponding normative expectation that he has for me to do so will ground a *complaint* against me when I decide to eat cold cereal and do philosophy instead.[106] If my failure to act on a relationship-dependent reason is to correspond to a complaint-grounding normative expectation of the other participant, that relationship-dependent reason must already be strong or decisive. There is some evidence that Scheffler agrees, for he also writes,

> [I]f the source of my reason to respond to your needs and desires lies in the value of our relationship, *and that reason is compelling*, then my reason for action is complemented by your entitlement to expect that I shall respond. The very same consideration that gives me reason to act on your behalf gives you reason to complain if I do not. (2010: 111, italics added)

If this is right, then it would be nice to have something to say about when a relationship-dependent reason comes to be sufficiently strong or compelling enough to have a corresponding complaint-grounding entitlement in the other person. It is not obvious how this should be done. It wouldn't work, for example, simply to measure the strength of a relationship-dependent reason in terms of the weightiness of the other's interests that are at stake. That my neighbor happens to have an extremely strong interest in eating from my grill today wouldn't suffice to ground his complaint.

So it may be more promising to look at the significance of the relationship. However, it seems that my deeply valuable relationship to my child might nevertheless sometimes generate very slight relationship-dependent reasons for me to act in certain ways on my child's behalf. Moreover, it is conceivable that a particular interpersonal relationship might generate strong relationship-dependent reasons even though the relationship itself (i.e., the token, not the type) is relatively trivial in value. The cost of

[106] Or, if this does give my neighbor a complaint against me, it won't have the sort of normative significance Scheffler seeks. The complaint won't correspond to any moral wrongdoing.

denying this possibility is great if we are supposing that the strength of relationship-dependent reasons is a function of the relative value of the relationship. While my associative duties towards my own child might be secure (it being a relatively high value relationship), any associative duties towards my colleagues, students, or neighbors would be suspect.

Now it might be thought that the strength of a relationship-dependent reason is a function of what is stake in the relationship itself. So if a failure to respect such a reason would destroy the relationship, this should strengthen the reason. One problem with this proposal is, as Scheffler himself observes, that "it is often wrong to neglect one's special responsibilities even when doing so would not destroy the relationship giving rise to those responsibilities" (2004: 265). Another problem is that whether or not the relationship will be damaged will often depend on whether the participants in the relationship *regard* the reason in question as providing a moral duty. This suggests that the decisive factor in determining the strength of a relationship-dependent reason might simply be what the participants, or at least one of them, take the significance of that reason to be. This may make the question of special duties counter-intuitively vulnerable to perversion by the unreasonable expectations of any participant in a valuable special relationship. Do I suddenly have a duty to call my friend on his birthday just because, strangely, he has unilaterally come to believe that this is a requirement of the friendship?

It seems rather that a given type of relationship creates reasons peculiar to its type, and that some of these reasons will be stronger than others. Whether or not a given reason of a given relationship is a particularly strong one will depend on the role of that reason within the type of relationship, and whether or not, in the given type of

circumstance, it would be acceptable to permit members of the relationship to refrain from acting on the reason. The reason that my student has a complaint against me if I neglect to read her paper hasn't directly to do with the relative significance of the student's interests in having her paper read, nor with the value of my relationship to this particular student. Rather, the legitimacy of the student's complaint is better explained by the unacceptability of a general permission for teaching fellows to neglect to read their students' papers in similar circumstances.[107] On my own view, it is unacceptability of this kind that generates special reasons of the sort that, with the addition of deontic direction, can become special duties owed to others.

I propose, then, that the valuable relationship that generates relationship-dependent reasons is not the particular history of interpersonal interactions that Scheffler has in mind. The valuable relationship is the type, not the token. It is this which helps to explain why I have a compelling enough, relationship-dependent reason to grade the student's paper even though it would not be very detrimental to the student or to anyone else if I neglected to do so, and even though my relationship with the student is not, of itself, highly valuable, at least as interpersonal relationships go. I say that my relationship-dependent reason to grade the paper is still *compelling enough* in the sense that it corresponds to a complaint she has against me if I fail to respect it.

In speaking, then, of relationship types as opposed to relationship tokens, I have in mind the sets of interlocking norms and normative expectations comprising these

[107] Raz proposes something similar for a broad class of rights, noting that "quite commonly the value of a right, the weight it is to be given or the stringency with which it is to be observed do not correspond to its value to the right-holder" (1992: 128). So Raz argues, "The protection of many of the most cherished civil and political rights in liberal democracy is justified by the fact that they serve the common or general good. Their importance to the common good, rather than their contribution to the well-being of the right-holder, justifies the high regard in which such rights are held…." (135). As examples, Raz points to freedoms of contract and expression. Elsewhere he points to the right of journalists to protect their sources (1988: 247-48). Raz also thinks, "Arguable the same is true of all rights" (1992: note 6).

relationship types. This puts me directly at odds with Scheffler's insistence that the relationships from which associative duties derive be understood as "ongoing bonds between individuals who have a shared history" (2010: 115). As I see it, Scheffler's way of understanding the object of value makes it difficult to explain adequately why relationship-dependent reasons sometimes become obligatory or non-optional, and why sometimes they don't. Scheffler, it seems, is unable to appeal to the larger societal practice in which the norm-defined relationships are respected. For the same reason, Scheffler's insistence would block my attempt to extend his account of associative duties to duties of rescue in the pond-like cases. In the paradigmatic pond-like cases the agents have no special shared history with one another.

Scheffler defends his insistence by claiming that unless relationships are understood in his historically robust sense, we may be incapable of appealing to the relationship to explain the relationship-dependent norms. The idea is that, unless the relationships have content beyond those norms of the relationship, the relationship cannot explain the norms (2010: 123).

Is this right? Suppose we define a hypothetical relationship that consists merely of a defined set of interlocking norms specifying the terms according to which participants should relate. This would not, for Scheffler, count as a proper relationship. However, of such a hypothetical, purely norm-defined relationship we might acknowledge this: if such a relationship *were* realized (i.e., if such norms were actually respected) such a relationship would be non-instrumentally valuable. Inspired by Scheffler, we might then say that such recognized value provides us with reason to realize the relationship. This simply means that we have reason to respect the set of

norms that constitutes the valuable relationship. Since this reason to respect the norms derives from the non-instrumental value of so doing—which is something that we are still yet to do—it must be that this reason does not derive from any historically realized relationship of the sort Scheffler insists upon. Insofar, then, as the recognizable prospective value of enacting this relationship explains our reasons for doing so, Scheffler must be wrong about the explanatory need for relationships in his historically robust sense.

Or, perhaps Scheffler's idea is this: if the relationship is nothing more than the set of norms defining it, then any judgment that such a relationship is valuable may be nothing more than the affirmation that the relevant norms are legitimate. Such an affirmation would not *explain* the legitimacy of the norms in question. It is a dubious claim, however, that the judgment about such a relationship's value is *nothing more* than a judgment that each of the specific norms constituting that relationship is legitimate. A relationship of this kind should be regarded as a larger whole of which the many particular norms comprising it are parts. We can, it seems to me, often explain the value of the parts by appealing to the more recognizable value of the whole. When considered apart from the whole, the particular norms comprising the whole may look unjustified, suspect, or even absurd. Analogously, if we lift out a particular clause from the larger policy proposal it is in, then the particular clause might look pointless. Likewise, a particular rule in a game may seem, in isolation, absurd. The same is true of the norms of interpersonal relationships. A norm requiring one person to avoid doing things that might provoke jealously in another may only be justified within the more complex relationship of two lovers.

Whether or not we should think of the assumed practice of mutual aid as a "relationship," there is further reason why we ought to be able to extend Scheffler's account so as to include that assumed practice as a source of special reasons. For Scheffler, there are membership-dependent reasons in addition to relationship-dependent reasons. As with relationship-dependent reasons, one may be "required or obligated to act on" membership-dependent reasons (2010: 112).

What is most noteworthy about this class of reasons is the way Scheffler characterizes the value from which they derive. He suggests that one's membership in a group or an association may be non-instrumentally valued "because of the bonds of trust and solidarity that members share" (107). This sounds very much like the value that Miller describes when characterizing the value of the assumed practice of mutual aid. Miller and Scheffler seem here to be talking about the same kind of value. Instead of membership-dependent reasons, we might even call them "practice-dependent reason."[108]

Here, then, is what I conclude about Scheffler's proposal about associative duties and the prospects for incorporating the duty of rescue in the pond-like cases among them. First, Scheffler has no good reason for resisting the derivation of special moral duties from values of the sort Miller describes, the value of a societal practice in which people are expected to rescue one another in pond-like cases. Second, Scheffler needs to improve his explanation for how membership and relationship-dependent reasons sometimes become obligatory. My proposal, in terms of unacceptability judgments, offers a way for Scheffler to do this. Third, Scheffler's proposal might be improved by interpreting relationships and memberships not as historical interaction between

[108] Scheffler acknowledges that three categories ("project-dependent," "relationship-dependent," and "membership-dependent") "may not be exhaustive" (2010: 112).

individuals, but rather as sets of interlocking norms whose value is realized when the norms are respected.

6.2 The category of special duties

The imperiled person in Shallow Pond is to me a complete stranger, toward whom I have made no special commitment. I have also not, at least in any obvious way, accepted the responsibilities of any special role or office requiring me to perform such rescues. It is not as if I am employed as a lifeguard. It appears as if performing the rescue in Shallow Pond is something that I owe to the imperiled person simply as another person, and because I happen to be there. As such, this is the sort of duty that Scheffler's characterization of special duties appears designed to preclude, since the imperiled person is apparently neither someone with whom I have had any "significant sort of interactions," nor someone to whom I stand in any "significant sort of relation." Or, if we insist that I do stand in a significant sort of relationship with this person simply in virtue of occupying a circumstance that places me under a duty to that person, then we arguably weaken Scheffler's characterization of special duties to the point of collapsing any meaningful distinction between special and general duties.

That there may be no important and fundamental distinction between special and general duties strikes me as plausible, especially if special duties are taken to include contractual duties, relationship-based duties, role-based duties, and membership-based duties. If the only thing that this broad class of duties is taken to have in common is that they apply only in certain, specially defined circumstances of one kind or another, then it may be difficult to find a non-special duty that is different in any important way.

I suspect, therefore, that it isn't worth our while to try to decide whether we should regard the duty to rescue in the pond-like cases as a special duty. It is better just to observe that such rescue duties have interesting similarities to many of the sorts of duties often classified as special duties (such duties often make positive demands and exhibit the deontic features). These similarities suggest that the category of special duties may have been drawn artificially narrow, rendering questionable the attempts to account for special duties in terms of the features encompassed by that category. If, as I have claimed, duties of rescue are, in their most theoretically troublesome features, a lot like prototypical special duties, then a successful general account of "special duties" might need to cover them both. Or, at least, in attempting to offer a general account of special duties, or a subset thereof, we would do well to avoid focusing entirely on the explanatory resources made available only through the presence of significant relationships or transactions. If, moreover, my account of the stringent duty of rescue in the pond-like cases is successful, it is the kind of proposal that could provide fresh insight. Moreover, because my proposal doesn't need to appeal to any of those elements most typically associated with special duties, it is the kind of account that might shed light on moral duties more generally, both special or otherwise.

6.3 Dismissing the over-demandingness loophole

I argued in chapter one that the duty to rescue in the pond-like cases poses a distinctive threat of over-demandingness insofar as such a duty appears to be general, positive, and perfect, thereby avoiding all the categories useful for limiting the demandingness of a duty ("special," "negative," or "imperfect"). The more we rely upon

such categories for limiting demandingness, the more we should be troubled by duties that fit none of these categories. We should then also be troubled by the puzzle of explaining the significance of these categories. Why should it be that a duty tends to be perfect if and only if it is negative or special? Why should it be that positive, general duties tend not to be perfect?

Such questions are less mysterious if we accept, as I have, that moral duties generally are sensitive to the cumulative costs to individuals of compliance. A reason that only negative or special duties tend to be perfect is that positive, general duties would tend to demand too much if they were also perfect. In accepting, moreover, that the significance of these categories is in part explained by their ability or inability to limit the cumulative costs to individuals of compliance, we can further accept that the negative nature of a duty, or the fact that the duty involves voluntarily assumed obligations or depends on special relationships, is not the fundamental reason why the duty in question is stringent or perfect. Rather, what fundamentally matters is that the costs to an individual of complying with such stringent duties are sufficiently low relative to the value of fulfilling them, individually and collectively. Under my proposal, this implies that we should not regard as deeply significant the category differences between, on the one hand, negative or special duties, and, on the other, the stringent duty of rescue in the pond-like cases. Of importance is that the costs of compliance with the duty to rescue in the pond-like cases are, as with these other duties, securely limited.

I am suggesting, therefore, that we should not take seriously the apparent loophole for over-demandingness in the standard taxonomy of moral duties that I discussed in chapter one. The stringent duty of rescue in the pond-like cases is not made problematic

by its status as general and positive. So long as the circumstances of the duty's application are sufficiently restricted, its general and positive character poses no distinctive problem. The greater task was to provide a principled basis for such restrictions.

7. Conclusions

The famine-relief argument is of interest because it holds the promise of yielding a self-standing problem of over-demandingness, one that might challenge widespread and basic assumptions about morality. I attempted to identify the central problem posed by that argument, and to defend, in the face of its challenge, a comfortably moderate position. To gain plausibility for this defense, I avoided making special exceptions for rescue duties, or duties of assistance, as a class. My aim was instead to make room for the moderate view of such duties, showing how our rescue duties, so interpreted, are unexceptional, exhibiting just those features that we plausibly encounter in moral duties more generally. This aim led me to two broad ideas. The first is that considerations of cumulative costs to an individual of compliance directly limit the scope and content of our moral duties generally. When such costs of compliance dramatically change, we shouldn't assume that the moral duties remain fixed. Surprising changes in circumstance can mean surprising changes in our duties. The second broad idea is a picture of how stringent duties can derive from the unacceptable consequences that a type of behavior, in aggregate, has or would have, regardless of the merely trivial consequences that the prohibited behavior may have in isolation. An effect of appealing to such broad ideas is

my proposal has implications spreading beyond my particular application. Such implications are, I believe, a further payoff, offering potentially helpful perspectives on the category of special duties, the status of negative duties (relative to positive ones), the features of deontological constraints, and consequence-based evaluations of actions whose individual effects are negligible.

The plausibility of my account of rescue duties depends in part on the availability and the comparative merits of other options. While I directly discussed only some of these alternatives, my clarifications of the basic challenges provide measures against which the alternative moderate proposals may be assessed. Instead of giving a better defense of the moderate position, some will prefer to reject the assumptions underlying that position. An obvious option here would be to reject those assumptions wherein I try to accommodate my own intuitive judgments about cases, judgments that may strike others as less compelling. Rather than agreeing with my conclusions about, in particular, Distant Rescue, Nightmare, and Intergalactic Killer,[109] some may feel that certain of the challenges I address are unnecessary. Some may find it best to join the efforts to further strengthen the debunking explanation of the moderate's unwillingness to accept certain costly requirements of rescue.

Despite the remaining possibilities, I take myself to have shown the reasonableness of the moderate position in the face of the general skeptical worries endemic to that position. It is a view of our rescue duties that someone can provisionally adopt, in good faith and without egregious moral insensitivity, despite awareness of the best currently available arguments on the matter.

[109] Pages 40, 86 and 126, respectively.

BIBLIOGRAPHY

Arneson, Richard J. 2004. Moral Limits on the Demands of Beneficence? In *The Ethics of Assistance.*, ed. Deen K. Chatterjee, 33-58. Cambridge: Cambridge University Press.

Arthur, John. 1977. Rights and the Duty to Bring Aid. In *World Hunger and Morality.*, eds. William Aiken, Hugh LaFollette. 2nd ed., 39-50 Prentice Hall.

———. 2002. Famine Relief and the Ideal Moral Code. In *Ethics for Everyday.*, ed. David Benatar. 1st ed., 520-536. New York: McGraw-Hill.

Ashford, Elizabeth. 2003. The Demandingness of Scanlon's Contractualism. *Ethics* 113: 273-304.

Badhwar, Neera K. 2006. International Aid: When Giving Becomes a Vice. *Social Philosophy and Policy* 23 (01): 69.

Cohen, Andrew I. 2005. Famine Relief and Human Virtue. In *Contemporary Debates in Applied Ethics.*, eds. Christopher Heath Wellman, Andrew I. Cohen, 326-342. Malden, MA: Blackwell Publishing.

Cohen, L. Jonathan. 1981. Who is Starving Whom? *Theoria* 47: 65-81.

Cullity, Garrett. 1994. International Aid and the Scope of Kindness. *Ethics* 105: 99-127.

———. 1995. Moral Character and the Iteration Problem. *Utilitas* 7: 289-99.

———. 1996. The Life-Saving Analogy. In *World Hunger and Morality.*, ed. Hugh LaFollette William Aiken. 2nd ed., 51-69. Englewood Cliffs, NJ: Prentice-Hall.

———. 2000. Pooled Beneficence. In *Imperceptible Harms.*, ed. Michael J. Almeida, 9-42. Dordrecht: Kluwer.

———. 2003. Asking Too Much. *The Monist* 86 (3): 402-18.

———. 2004. *The Moral Demands of Affluence*. Oxford; New York: Clarendon Press; Oxford University Press.

———. 2009. Demandingness and Arguments from Presuppositions. In *The Problem of Moral Demandingness.*, ed. Timothy Chappell, 8-34. New York: Palgrave Macmillan.

Dower, Nigel. 2003. World Hunger. In *Oxford Handbook of Practical Ethics.*, ed. Hugh LaFollette, 643-669. New York: Oxford University Press.

Fabre, Cécile. 2006. *Whose Body Is It Anyway?: Justice and The Integrity of The Person*. Oxford; New York: Clarendon Press; Oxford University Press.

Feinberg, Joel. 1984. The Moral and Legal Responsibilities of the Bad Samaritan. *Criminal Justice Ethics* 3 (1): 56-69.

———. 1987. Failures to Prevent Harm. In *The Moral Limits to The Criminal Law, vol. 1: Harm to Others*, 126-186. Oxford University Press.

Fishkin, James S. 1982. *The Limits of Obligation*. New Haven, Conn.: Yale University Press.

Foot, Philippa. 1972. Morality as a System of Hypothetical Imperatives. *The Philosophical Review* 81 (3): 305-16.

Foot, Philippa. 1985. Utilitarianism and the Virtues. *Mind* 94 (374): 196-209.

Gewirth, Alan. 1987. Private Philanthropy and Positive Rights. *Social Philosophy and Policy* 4: 55-78.

Gomberg, Paul. 2002. The Fallacy of Philanthropy. *Canadian Journal of Philosophy* 32: 29-66.

Goodin, Robert. 1985. *Protecting the Vulnerable: A Reanalysis of Our Social Responsibilities*. Chicago: University of Chicago Press.

Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.

Hare, Caspar. forthcoming. Obligations to Merely Statistical People (September 2012 draft), *Journal of Philosophy*. http://web.mit.edu/~casparh/www/Papers/CJHareObligationstoStatistical.pdf.

Hare, R. M. 1965. *Freedom and Reason*. Oxford: Oxford University Press.

Hart, H. L. A. 1955. Are There Any Natural Rights? *The Philosophical Review* 64 (2) (Apr.): 175-91.

Herman, Barbara. 1984. Mutual Aid and Respect for Persons. *Ethics* 94: 577-602.

———. 2000. Morality and Everyday Life. *Proceedings and Addresses of the American Philosophical Association* 74 (2): 29-45.

———. 2001. The Scope of Moral Requirement. *Philosophy and Public Affairs* 30 (3): 227-56.

Hershenov, David B. 2002. A Puzzle About the Demands of Morality. *Philosophical Studies* 107 (3): 275-90.

Hill, Thomas E., Jr. 1992. How a Kantian Might Think about Terrorism. In *Dignity and Practical Reason.* Ithaca, New York: Cornell University Press.

———.1997. Reasonable Self-Interest. *Social Philosophy and Policy* 14 (01): 52.

Hooker, Brad. 1999. Review: Sacrificing for the Good of Strangers-Repeatedly. *Philosophy and Phenomenological Research* 59 (1) (Mar.):177-181.

———. 2002. *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Oxford University Press.

Igneski, Violetta. 2006. Perfect and Imperfect Duties. *Social Theory and Practice* 32 (3).

James, Susan. 1982. The Duty to Relieve Suffering. *Ethics* 93 (1): 4-21.

Jamieson, Dale. 2005. Duties to the Distant: Aid, Assistance, and Intervention in the Developing World. *The Journal of Ethics* 9 (1/2, Current Debates in Global Justice): 151-170.

Kagan, Shelly. 1984. Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation. *Philosophy and Public Affairs* 13 (3): 239-54.

———. 1988. The Additive Fallacy. *Ethics* 99 (1) (Oct.): 5-31.

———. 1989. *The Limits of Morality*. Oxford Ethics Series. Oxford; New York: Clarendon Press; Oxford University Press.

———. 1992. The Structure of Normative Ethics. *Philosophical Perspectives* 6 (Ethics): 223-42.

Kamm, F. M. 2001. *Morality, Mortality Volume II: Rights, Duties, and Status*. New York: Oxford University Press.

———. 2007. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford Ethics Series. Oxford; New York: Oxford University Press.

Kekes, John. 2002. On the Supposed Obligation to Relieve Famine. *Philosophy* 77: 503-17.

Lewis, David K. 2000. Illusory Innocence. In *Papers in Ethics and Social Philosophy: Volume 3.*, 152-158. Cambridge University Press.

Lewis, Clarence Irving. 1946. *An Analysis of Knowledge and Valuation*. [1st] ed. Vol. 7th ser., 1945. La Salle, Ill.: The Open court publishing company.

Lomasky, Loren E. 1995. Justice to Charity. *Social Philosophy and Policy* 12 (02): 32.

McGinn, Colin. 1996. Saints Elsewhere. *New Republic.* October 14.

McKinsey, Michael. 1981. Obligations to The Starving. *Nous* 15 (3): 309-23.

McMahan, Jeff. 2000. Moral Intuition. In *The Blackwell Guide to Ethical Theory.*, ed. Hugh LaFollette, 92-110. Malden, MA: Blackwell.

Mill, John Stuart. 1993. *Utilitarianism*. In *British Philosophy: 1600-1900.*, ed. Mark C. Rooks. Charlottesville: InteLex Corporation.

Miller, Richard W. 2004. Beneficence, Duty and Distance. *Philosophy and Public Affairs* 32 (4): 357-83.

Mothersill, Mary. 1993. The Limits of Morality, by Shelly Kagan. *The Journal of Philosophy* 90 (10): 537-44.

Mulgan, Tim. 1997. Two Conceptions of Benevolence. *Philosophy & Public Affairs* 26 (1) (Winter): 62-79.

Murphy, Liam B. 1993. The Demands of Beneficence. *Philosophy and Public Affairs* 22 (4): 267-92.

———. 1997. A Relatively Plausible Principle of Beneficence: Reply to Mulgan. *Philosophy & Public Affairs* 26 (1) (Winter): 80-86.

———. 1998. Institutions and the Demands of Justice. *Philosophy & Public Affairs* 27 (4) (Autumn): 251-291.

———. 2000. *Moral Demands in Nonideal Theory*. Oxford Ethics Series. New York: Oxford University Press.

Nagel, Thomas. 1978. *The Possibility of Altruism*. Princeton: Princeton University Press.

———. 1986. *The View From Nowhere*. New York: Oxford University Press.

———. 1995. *Equality and Partiality*. New York: Oxford University Press.

Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.

Odegard, Douglas. 1989. Charity and Moral Imperatives. *Theoria* 55: 81-94.

O'Neill, Onora. 1986a. *Faces of Hunger: An Essay on Poverty, Justice, and Development*. Studies in Applied Philosophy. Vol. 3. London ; Boston: G. Allen & Unwin.

———. 1986b. The Moral Perplexities of Famine and World Hunger. In *Matters of Life and Death.*, ed. Tom Regan. 2nd ed., 294-337. New York: Random House, Inc.

Parfit, Derek. 1989. *Reasons and Persons*. Reprint with corrections. Oxford: Clarendon Press.

Prichard, H. A. 1912. Does Moral Philosophy Rest on a Mistake? *Mind* 21 (81) (Jan.): 21-37

Pufendorf, Samuel. 1717. *Of the Law of Nature and Nations. Eight Books.* Trans. Basil Kennet. London: printed for R. Sare, R. Bonwicke, T. Goodwyn, J. Walthoe, M. Wotton, S. Manship, R. Wilkin, B. Tooke, R. Smith, T. Ward, and W. Churchill.

Quinn, Warren S. 1990. The Puzzle of the Self-Torturer. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 59 (1) (May): 79-90.

Rachels, James. 1975. Active and Passive Euthanasia. *New England Journal of Medicine* 292 (2): 78-80.

———. 1979. Killing and Starving to Death. *Philosophy* 54: 154-66.

Rawls, John. 1993. *Political Liberalism*. John Dewey Essays in Philosophy. Vol. 4. New York: Columbia University Press.

———. 1999a. *A Theory of Justice*. Rev ed. Cambridge, Mass.: Belknap Press of Harvard University Press.

———. 1999b. *The Law of Peoples: with "the Idea of Public Reason Revisited"*. Cambridge, Mass.: Harvard University Press.

Raz, Joseph. 1988. *The Morality of Freedom*. Oxford: Oxford University Press.

———. 1992. Rights and Individual Well-Being. *Ratio Juris* 5 (2): 127-42.

Ripstein, Arthur. 2000. Three Duties to Rescue: Moral, Civil, and Criminal. *Law and Philosophy* 19 (6, The Moral and Legal Limits of Samaritan Duties) (Nov.): 751-779.

Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, Mass.: Belknap Press of Harvard University Press.

———. 2003a. *The Difficulty of Tolerance: Essays in Political Philosophy*. Cambridge; New York: Cambridge University Press.

———. 2003b. Replies. *Ratio* xvi (4): 424-39.

———. 2003c. Review: Reply to Gauthier and Gibbard. *Philosophy and Phenomenological Research* 66 (1): 176-89.

———. 2007. Wrongness and Reasons: A Reexamination. In *Oxford Studies in Metaethics.*, ed. Russ Shafer-Landau. Vol. 2, 5. New York: Oxford University Press.

Scheffler, Samuel. 1986. Morality's Demands and Their Limits. *The Journal of Philosophy* 83 (10): 531-7.

———. 1992. *Human Morality*. New York: Oxford University Press.

———. 1994. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Rev ed. Oxford; New York: Clarendon Press; Oxford University Press.

———. 1997. Relationships and Responsibilities. *Philosophy & Public Affairs* 26 (3) (Summer): 189-209.

———. 2003. *Boundaries and Allegiances*. Oxford Scholarship Online edition. Oxford Scholarship Online.

———. 2004. Projects, Relationships, and Reasons. In *Reason and Value: Themes From the Moral Philosophy of Joseph Raz.*, eds. R. Jay Wallace, Philip Pettit, Samuel Scheffler and Michael Smith, 247-269. New York: Oxford University Press.

———. 2005. The Division of Moral Labour. *Proceedings of the Aristotelian Society, Supplementary Volumes* 79: 229-253.

———. 2010. Morality and Reasonable Partiality. In *Partiality and Impartiality.*, ed. Brian Feltham, John Cottingham, 98-130. New York: Oxford University Press.

Schmidtz, David. 2000. Islands in a Sea of Obligation: Limits of the Duty to Rescue. *Law and Philosophy* 19: 683-705.

Schroeder, S. Andrew. forthcoming. Imperfect Duties, Group Obligations, and Beneficence. *Journal of Moral Philosophy*. http://www.cmc.edu/pages/faculty/ASchroeder/docs/SASImperfectDuties.pdf

Sidgwick, Henry. 1981. *The Methods of Ethics*. 7th ed. Indianapolis: Hackett Pub. Co.

Singer, Peter. 1972. Famine, Affluence and Morality. *Philosophy and Public Affairs* 1 (3): 229-43.

———. 1974. Sidgwick and Reflective Equilibrium. *The Monist* 58: 490-517.

———. 1977. Reconsidering the Famine Relief Argument. In *Food Policy: The Responsibility of the United States in the Life and Death Choices.*, eds. Peter G. Brown, Henry Shue, 36-53. New York: Free Press.

———. 1988. Reasoning Towards Utilitarianism. In *Hare and Critics.*, eds. D. Seanor, N. Fotion, 147-159. Oxford: Clarendon Press.

———. 1995. *How Are We to Live?: Ethics in an Age of Self-Interest*. Amherst, N.Y.: Prometheus Books.

———. 1999a. A Response. In *Singer and His Critics.,* ed. Dale Jamieson, 269-335. Malden, MA: Blackwell.

———. 1999b. Review of Unger's *Living High and Letting Die*. *Philosophy and Phenomenological Research* 59 (1) (Mar.): 183-187.

———. 2004. Outsiders: Our Obligations to Those Beyond Our Borders. In *The Ethics of Assistance.*, ed. Deen K. Chatterjee, 11-32. Cambridge: Cambridge University Press.

———. 2007. Review Essay on *The Moral Demands of Affluence*. *Philosophy and Phenomenological Research* LXXV (2): 475-83.

———. 2009. *The Life You Can Save: Acting Now To End World Poverty*. 1st ed. New York: Random House.

———. 2011. *Practical Ethics*. 3rd ed. New York: Cambridge University Press.

Singer, Peter, and Dale Jamieson. 1999. *Singer and His Critics*. Philosophers and Their Critics. Vol. 8. Oxford, UK; Malden, Mass.: Blackwell Publishers.

Slote, Michael. 1977. The Morality of Wealth. In *World Hunger and Moral Obligation.*, eds. William Aiken, Hugh LaFollette, 124-147. Englewood Cliffs, NJ: Prentice Hall.

Smith, Patricia. 1990. The Duty to Rescue and the Slippery Slope Problem. *Social Theory & Practice* 16 (1) (Spring): 19-41.

Sobel, David. 2007. The Impotence of the Demandingness Objection. *Philosophers' Imprint* 7 (8).

Sreenivasan, Gopal. 2010. Duties and Their Direction. *Ethics* 120 (3) (April): 465-494.

Tännsjö, Torbjörn. 2000. Must We, the Reader and I, Make Great Sacrifices in Order to Save Starving Children? *Theoria* 66 (1): 115-22.

Temkin, Larry S. 1986. A Continuum Argument for Intransitivity. *Philosophy and Public Affairs*. 25 (3) (Summer): 175-210.

Thomas, Alan. 2009. Consequentialism, Integrity and Demandingness. In *The Problem of Moral Demandingness: New Philosophical Essays.*, ed. Timothy Chappell, 123-147. New York: Palgrave Macmillan.

Thomson, Judith Jarvis. 1971. A Defense of Abortion. *Philosophy & Public Affairs* 1 (1) (Autumn): 47-66.

———. 1990. *The Realm of Rights*. Cambridge, Mass.: Harvard University Press.

Trammell, Richard L. 1975. Saving Life and Taking Life. *The Journal of Philosophy* 72 (5) (Mar. 13): 131-7.

Unger, Peter K. 1996. *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.

Wallace, R. Jay. unpublished. The Deontic Structure of Morality, http://philosophy.berkeley.edu/file/2/Deonticstructure-final.pdf (accessed February 1st, 2013).

Walzer, Michael. 1983. *Spheres of Justice: A Defense of Pluralism and Equality*. New York: Basic Books.

Whelan, John M. 1991. Famine and Charity. *Southern Journal of Philosophy* 29: 149-66.

Williams, Bernard. 1973. A Critique of Utilitarianism. In *Utilitarianism, For and Against.*, 77-150. Cambridge: Cambridge University Press.

———. 1981. Persons, Character, and Morality. In *Moral Luck*, 1-19. Cambridge: Cambridge University Press.